arXiv:1509.08362v1 [math.ST] 28 Sep 2015

Blocking Strategies and Stability of Particle Gibbs Samplers

Sumeetpal S. Singh Fredrik Lindsten Eric Moulines

September 28, 2015

Abstract

Sampling from the conditional (or posterior) probability distribution of the latent states of a Hidden Markov Model, given the realization of the observed process, is a non-trivial problem in the context of Markov Chain Monte Carlo. To do this Andrieu et al. (2010) constructed a Markov kernel which leaves this conditional distribution invariant using a Particle Filter. From a practitioner's point of view, this Markov kernel attempts to mimic the act of sampling all the latent state variables as one block from the posterior distribution but for models where exact simulation is not possible. There are some recent theoretical results that establish the uniform ergodicity of this Markov kernel and that the mixing rate does not diminish provided the number of particles grows at least linearly with the number of latent states in the posterior. This gives rise to a cost, per application of the kernel, that is quadratic in the number of latent states which could be prohibitive for long observation sequences. We seek to answer an obvious but important question: is there a different implementation with a cost per-iteration that grows linearly with the number of latent states, but which is still stable in the sense that its mixing rate does not deteriorate? We address this problem using blocking strategies, which are easily parallelizable, and prove stability of the resulting sampler.

1 Introduction

Let $\{(X_t, Y_t) \in (\mathsf{X}, \mathsf{Y}), t \in \mathbb{N}_+\}$ be a Markov chain evolving according to a (general state space) hidden Markov model (HMM). That is, the *state process* $\{X_t,\ t \in \mathbb{N}_+\}$ is a Markov chain with *state space* X, evolving according to a Markov transition kernel M and with initial distribution μ (i.e., μ is the distribution of the state at time $t = 1$). The sequence $\{X_t,\ t \in \mathbb{N}_+\}$ is not directly observed and inference needs to be carried out based on the observations $\{Y_t,\ t \in \mathbb{N}_+\}$ only. Conditionally on $\{X_t,\ t \in \mathbb{N}_+\}$, the observations $\{Y_t,\ t \in \mathbb{N}_+\}$ are independent. The conditional density of Y_t given $X_t = x_t$, with respect to some dominating measure is denoted by $g(x_t, y_t)$.

We will work under the assumption that a *fixed* sequence of observations $y_{1:n} := (y_1, \ldots, y_n)$ is available, where n is some final time point. The key object of interest is then the *joint smoothing distribution* (JSD) $\phi(\mathrm{d}x_{1:n})$, which is the probability distribution of $X_{1:n} := (X_1, \ldots, X_n)$ conditioned on $Y_{1:n} = y_{1:n}$. Markov chain Monte Carlo can be used to simulate from the JSD, for example the Gibbs scheme that uses Metropolis-Hastings samplers to update the state variables $\{X_t\}_{t=1}^n$ one-at-a-time, with all the other variables kept fixed (Carter and Kohn, 1994; Frühwirth-Schnatter, 1994). However, the typically strong dependencies between consecutive states in the state sequence can cause this method to mix very slowly. As a result, this solution is often deemed as inefficient.

However, recent developments in *sequential Monte Carlo* (SMC) methods have had a significant impact on the practise of MCMC. The SMC methodology—which combines sequential importance sampling and resampling— has long been established as a key technique for

approximating the JSD in general HMMs, see e.g., Doucet et al. (2000); Del Moral (2004); Doucet and Johansen (2011) for introductions, applications, and theoretical results. In a seminal paper by Andrieu et al. (2010), this key strength of SMC was exploited to construct effective (SMC based) high-dimensional MCMC kernels, resulting in so called particle MCMC (PMCMC) methods.

We will consider specifically an instance of PMCMC referred to as *particle Gibbs* (PG) Andrieu et al. (2010) (see also Chopin and Singh (2015)). The PG algorithm of Andrieu et al. (2010) defines (via an SMC construction) a Markov kernel which leaves the full JSD invariant. From a practitioner's point of view, PG can be regarded as a technique to mimic the behaviour of simulating all the state variables $X_{1:n}$ as one block from the JSD for models where exact simulation (from the JSD) is not possible.

There are some recent theoretical results for the PG (Markov) kernel $\mathcal{P}$ to back-up its good observed performance. Chopin and Singh (2015) show it is uniformly ergodic but do not shed light on the dependence of the rate on time n and particle number N. The PG kernel has been shown to be minorised by the JSD, i.e. $\mathcal{P} \geq \text{const.} \times \phi$, under weak conditions (Lindsten et al., 2015; Andrieu et al., 2015) (thus also implying uniform ergodicity). Under stronger (but common) forgetting conditions on the HMM, an explicit lower bound for the minorising constant, $(1 - \frac{1}{c(N-1)+1})^n$, has also been established in these works; here, $1 \geq c > 0$ is a model-dependent constant, N is the number of particles (of the SMC sampler) used to contruct the PG kernel and n is the number of observations. (See Kuhlenschmidt (2014) for an alternative proof of this rate, and a central limit theorem, using more routine SMC analysis techniques.) While the method is indeed uniformly geometrically ergodic for any $N \geq 2$, the explicit rate reveals that the convergence deteriorates as n increases if N is kept fixed. This effect, which is also clearly visible in practice (Lindsten and Schön, 2013; Chopin and Singh, 2015), is related to the well known path-degeneracy issue of SMC samplers, see e.g., Doucet and Johansen (2011). For the rate not to deteriorate, N must increase linearly with n, giving rise to an algorithm that costs n^2 per-iteration (or application) of the PG kernel, which may be impractical when n is large. (Note that a PG kernel implemented with N particles has a cost per-iteration of Nn.)

We seek to answer an obvious but important question: is there a different implementation of the PG kernel with a cost per-iteration that grows linearly with n, but which is still stable in the sense that its convergence rate does not deteriorate with n? Secondly, and equally important, is the issue of parallelizing the PG kernel to save on wall-clock, or execution, time. We provide an affirmative answer to the posed question and we use blocking to achieve both aims.

1.1 Summary of main results and related work

In the literature PG is often presented (and invoked) as an exact approximation of a *fully blocked* sampler and it is contrasted with a single-state Gibbs sampler. However, there is an intermediate route in-between these two extremes, namely to use PG as a component of a (partially) blocked Gibbs sampler. (This possibility was indeed pointed as a potential extension in the seminal PMCMC paper (Andrieu et al., 2010, p. 294).) As we will demonstrate in this paper, using PG within a blocked Gibbs sampler can result in very efficient samplers. We now give a *simplified* interpretation of our results.

In light of the convergence properties of the PG kernel discussed above, the convergence rate of the blocked PG kernel will now depend on the block size L and not on the total number of observations n. Thus the particle approximation of the ideal blocked Gibbs sampler does not deteriorate with total number of observations n provided the block sizes themselves do not increase with n. The main insight which we exploit next is that blocking can also be used to control the convergence properties of the ideal Gibbs sampler for an HMM. Specifically, we

show that under certain forgetting properties of the HMM it is possible to select a blocking scheme, using *overlapping* blocks, which results in a uniform rate of convergence for the ideal Gibbs sampler. Importantly, we show that that the rate of convergence after k complete sweeps is (see Theorems 4, 6)

$$|\phi\varphi - \mu\mathcal{P}^k_{\text{Ideal}}\varphi| \leq \lambda^k_{\text{Ideal}} \sum_{i=1}^{n} \text{osc}_i(\varphi)$$

where $\mathcal{P}_{\text{Ideal}}$ is the Markov kernel defined by one complete sweep and the rate λ_{Ideal} improves with increasing the overlap between blocks but is independent of n. Furthermore, we translate this result to the (more practical) case in which exact sampling from the block-conditionals is not possible and PG kernels are used to simulate from each of these conditionals instead. In this case the rate becomes (see Theorem 8)

$$\lambda_{\text{PG}} = \lambda_{\text{Ideal}} + \text{const.} \times \epsilon_N$$

where $0 \leq \epsilon_N \leq 1$ quantifies the effect of departing from the ideal block sampler when using the PG kernels instead. Specifically, $\epsilon_N \downarrow 0$ as the particle number N increases, i.e. as the PG kernel better approximates the ideal block sampler, but $\epsilon_N \uparrow 1$ as the size of the blocks increase. We analyse different blocking schemes (Theorem 8) and discuss in particular how the blocking can be selected to open up for straightforward parallelization of the sampler.

Our analysis is based on Wasserstein estimates, see e.g. recent works by Wang and Wu (2014); Rebeschini and van Handel (2014). Wang and Wu (2014) study the convergence properties of a Gibbs sampler in high dimensions under a Dobrushin condition. Our work is in the same vein, but we are in particular interested in *verifying* a related condition for the case of a blocked Gibbs sampler for an HMM. Furthermore, we study the convergence of the non-ideal blocked Particle Gibbs sampler, for which the results by Wang and Wu (2014) do not apply. Rebeschini and van Handel (2014) generalise the Dobrushin comparison theorem and consider applications in high-dimensional filtering.

It should be noted that refined PG algorithms, incorporating explicit updates of the particle ancestry either as part of the forward SMC recursion (Lindsten et al., 2014) or in a separate backward recursion (Whiteley, 2010; Whiteley et al., 2010), have been developed. These modified PG samplers have empirically been shown to work well with small number of particles and to be largely robust to n. Nevertheless, to date, no theoretical guarantee for the stability of these algorithms has been given. For specific blocking schemes, blocked PG is straightforwardly parallelizable. It is worth noting that we can replace our blocked PG kernels with the corresponding PG kernels that incorporate updates of the particle ancestry as discussed above.

The rest of this paper is organised as follows. Section 2 presents the HMM, the definition of the JSD, the blocking schemes and particle implementation of the ideal blocked Gibbs sampler. Section 3 presents the main results on the uniform ergodicty and rates of convergence of the various samplers. Section 4 presents the proof of the stability of the blocked PG sampler. The proofs of supporting technical results are given in the appendices.

2 Block Sampling and Particle Gibbs

2.1 Hidden Markov models

Let set X be the state space of the time- and space-homogeneous HMM. We will work under the assumption that a fixed sequence of observations $y_{1:n} := (y_1, \ldots, y_n)$ is available, where n is some final time point. The key object of interest is then the JSD, that is the probability

distribution of $X_{1:n}$ conditioned on $Y_{1:n} = y_{1:n}$. This will be the target distribution for the various Monte Carlo sampling algorithms that we shall consider throughout this work. With,

$$p_\mu(y_{1:n}) := \int \mu(\mathrm{d}x_1) g(x_1, y_1) \prod_{t=2}^{n} M(x_{t-1}, \mathrm{d}x_t) g(x_t, y_t) \ ,$$

denoting the density of the observations $Y_{1:n}$ (with respect to some dominating measure) we can write the density of the JSD as

$$\phi(x_{1:n}) := \frac{1}{p_\mu(y_{1:n})} \mu(x_1) g(x_1, y_1) \prod_{t=2}^{n} m(x_{t-1}, x_t) g(x_t, y_t) \ , \tag{1}$$

where, by abuse of notation, we use the same symbol ϕ both for the JSD and for its density. Note that $m(x_{t-1}, x_t)$ is the density of $M(x_{t-1}, \mathrm{d}x_t)$ with respect to some σ-finite dominating measure on X. For notational simplicity, we do not make explicit the dependence of ϕ on the fixed observation sequence in the notation in (1).

We will analyse the stability of our samplers under the following set of strong, but standard, mixing assumptions (Del Moral, 2004; Lindsten et al., 2015; Andrieu et al., 2015)

(S-1) There exists a probability measure ν on X such that $m(x_{t-1}, x_t)$ is the density of $M(x_{t-1}, \mathrm{d}x_t)$ with respect to ν. Furthermore, there exist positive constants σ_- and σ_+ and an integer $h \in \mathbb{N}$ such that

- (a) $m(x, x') \leq \sigma_+$ for all $x, x' \in \mathsf{X}$,
- (b) $\int \nu(dx_2) \dots \nu(dx_h) \left(\prod_{j=1}^{h} m(x_j, x_{j+1}) \right) \geq \sigma_-$ for all $x_1, x_{h+1} \in \mathsf{X}$.

(S-2) There exists a constant $\delta \geq 1$ such that for all $y \in \mathsf{Y}$, $\sup_x g(x, y) \leq \delta^{1/h} \inf_x g(x, y)$.

2.2 Block sampling

Before giving the algorithmic statements for the blocked Gibbs samplers that are analysed in this article, we introduce some notation that will be frequently used in the sequel. Let $I := \{1, \dots, n\}$ be the "index set" of the latent variables $X_{1:n}$. Let X^n be the n-fold Cartesian product of the set X and let the tuple $x = (x_i : i \in I) \in \mathsf{X}^n$. For $J \subset I$, we then write $x_J := (x_i : i \in J) \in \mathsf{X}^{|J|}$ (the restricted tuple). We also write x_{-i} as a shorthand for $x_{I \setminus \{i\}}$. The complement of J in I is denoted by $J^c := I \setminus J$. Given $y = (y_i : i \in J) \in \mathsf{X}^{|J|}$ and $z = (z_i : i \in J^c) \in \mathsf{X}^{|J^c|}$, we define $x = \{y, z\} = \{z, y\} \in \mathsf{X}^n$ to be the tuple such that $x_J = y$ and $x_{J^c} = z$.

Let ϕ_x^J denote (a version of) the regular conditional distribution of the variables X_J conditionally on $X_{J^c} = x_{J^c}$ under ϕ in (1). The Markov property of the HMM implies that ϕ_x^J depends on x only through the *boundary points* $x_{\partial J}$ where ∂J denotes the set of indices which constitutes the boundary of the set J

$$\partial J := \{t \in J^c : t + 1 \in J \text{ or } t - 1 \in J\} \ .$$

The Gibbs sampler generates samples from the JSD ϕ by iteratively sampling from its conditional distributions. Let $\mathcal{J} := \{J_1, \dots, J_m\}$ be a cover of I. A *blocked*, deterministic scan Gibbs sampler proceeds by sampling from the block-conditionals ϕ_x^J, $J \in \mathcal{J}$, in turn in some prespecified order. More precisely, assuming that we apply the blocks in the order J_1, J_2, etc.

and if the initial configuration of the sampler is given by $x' \in \mathsf{X}^n$, then the Gibbs sampler output is the Markov chain $\{X[k],\, k \in \mathbb{N}\}$ with $X[0] = x'$ and, given

$$X[lm + k - 1] = x \in \mathsf{X}^n, \quad l \in \mathbb{N},\, k \in \{1, \ldots, m\}\,, \tag{2}$$

$X[lm + k] = X'$ where $X'_{J_k^c} = x_{J_k^c}$ and we simulate $X'_{J_k} \sim \phi_x^{J_k}(\cdot)$. More generally, we can simulate X'_J from some kernel $Q^J(x, \cdot)$ with the conditional distribution ϕ_x^J as its invariant measure. That is, for any $x \in \mathsf{X}^n$

$$\int \phi_x^J(\mathrm{d}x'_J) Q^J(\{x_{J^c}, x'_J\}, A) = \phi_x^J(A), \qquad \text{for all measurable } A \subset \mathsf{X}^{|J|}\,. \tag{3}$$

Since ϕ_x^J depends on x only through the boundary points $x_{\partial J}$, it is natural to assume that $Q^J(x, \cdot)$ depends on x only through $x_{J \cup \partial J}$. For notational simplicity, we define

$$J^+ := J \cup \partial J$$

for any subset $J \subset I$. Thus for any $(x, z) \in \mathsf{X}^n \times \mathsf{X}^n$ we have, $Q^J(x, A) = Q^J(\{x_{J^+}, z_{I \setminus J^+}\}, A)$. We write $Q^J(x_{J^+}, \mathrm{d}x'_J)$ in place of $Q^J(x, \mathrm{d}x'_J)$ when wishing to to emphasize the dependence of the kernel on the components in J^+ of current configuration x. When we have $Q^J(x, \mathrm{d}x'_J) = \phi_x^J(\mathrm{d}x'_J)$ for every $J \in \mathcal{J}$, we refer to the sampler as an *ideal Gibbs sampler*. It follows that the Markov kernel corresponding to updating the block J is given by,

$$P^J(x, \mathrm{d}x') := \begin{cases} \phi_x^J(\mathrm{d}x'_J) \times \delta_{x_{J^c}}(\mathrm{d}x'_{J^c})\,, & \text{for the ideal Gibbs kernel,} \quad (4) \\ Q^J(x_{J^+}, \mathrm{d}x'_J) \times \delta_{x_{J^c}}(\mathrm{d}x'_{J^c})\,, & \text{for the non-ideal Gibbs kernel.} \quad (5) \end{cases}$$

Formally, the subsets in $\mathcal{J}$ can be an arbitrary cover of I. However, there are certain blocking schemes that are likely to be of most practical interest and these schemes therefore deserve some extra attention. To exemplify this, we consider the following restrictions on the blocks in $\mathcal{J}$.

(B-1) Each $J \in \mathcal{J}$ is an interval, i.e., $J = \{s, \ldots, u\}$ for some $1 \le s \le u \le n$. Furthermore, the blocks $J_1, \ldots, J_m$ are ordered in the sense that, for any $1 \le j < k \le m$, $\min(J_j) < \min(J_k)$ and $\max(J_j) < \max(J_k)$.

(B-2) Consecutive blocks may overlap, but non-consecutive blocks do not overlap and are separated. That is, for $1 \le j < k \le m$ with $k - j \ge 2$, $\max(J_j) < \min(J_k) - 1$.

In addition to ordering the blocks according to their minimum element, (B-1) avoids the case where one block is a strict subset of some other block. An illustration of (B-2), which requires that J_{k-1} and J_{k+1} do not cover J_k, is shown in Figure 1.

The Gibbs sampling scheme that (perhaps) first comes to mind is a systematic sweep from left to right. Under (B-1), this implies that the kernel corresponding to one *complete sweep* of the Gibbs sampler is given by:

L-R Assume (B-1). The Left-to-Right Gibbs kernel is given by $\mathcal{P} := P^{J_1} \cdots P^{J_m}$.

Assuming both (B-1) and (B-2), another blocking scheme that is of practical importance is the one that updates all the odd-numbered blocks first, then all the even-numbered blocks (or the other way around). We refer to this scheme as the *Parallel* (PAR) Gibbs sampler:

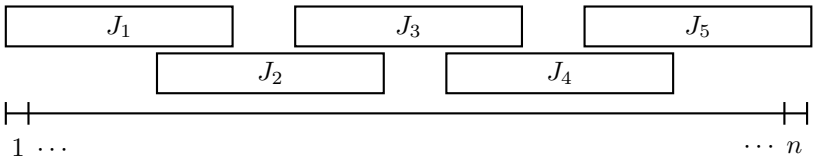

Figure 1: Example block configuration.

PAR Assume (B-1)–(B-2). The Parallel Gibbs kernel is given by $\mathcal{P} := \mathcal{P}_{\text{odd}}\mathcal{P}_{\text{even}}$ where

$$\begin{cases} \mathcal{P}_{\text{odd}} := P^{J_1} P^{J_3} \cdots P^{J_m}, \\ \mathcal{P}_{\text{even}} := P^{J_2} P^{J_4} \cdots P^{J_{m-1}}, \end{cases} \quad \text{or} \quad \begin{cases} \mathcal{P}_{\text{odd}} := P^{J_1} P^{J_3} \cdots P^{J_{m-1}}, \\ \mathcal{P}_{\text{even}} := P^{J_2} P^{J_4} \cdots P^{J_m}, \end{cases}$$

for odd/even m, respectively.

The reason for why we call this sampler *parallel* is that, under (B-1)-(B-2), it is possible to update all the odd blocks in parallel, followed by a parallel update of all the even blocks. This is because two consecutive odd (or even) blocks are separated by at least one element in I. Hence, if we have a total of $C/2$ processing units, we can simulate from both $\mathcal{P}_{\text{odd}}$ and $\mathcal{P}_{\text{even}}$ (and thus from $\mathcal{P}$) in constant time for any number of blocks $m \leq C$. Figure 1 shows an example of a blocking configuration which satisfies (B-1)-(B-2). This is the typical scenario that we have in mind for the PAR sampler.

Reversibility in MCMC is an important consideration. In Section 2.3 we use PG to define the kernels $Q^J(x_{J^+}, \mathrm{d}x'_J)$ for each J and the PG kernel is known to be reversible (Chopin and Singh, 2015). As such it is simple to define reversible block samplers using the following simple fact.

Lemma 1. *Let $\mathcal{J} = \{J_1, \ldots, J_m\}$ be an arbitrary cover of I and, for each $J \in \mathcal{J}$, let $P^J(x, \mathrm{d}x') = Q^J(x_{J^+}, \mathrm{d}x'_J) \times \delta_{x_{J^c}}(\mathrm{d}x'_{J^c})$ be the (possibly non-ideal) Gibbs kernel updating block J only. Assume that Q^J is reversible with respect to to ϕ_x^J for all $J \in \mathcal{J}$. Then,*

$$\phi(\mathrm{d}x) P^{J_1} \cdots P^{J_m}(x, \mathrm{d}x') = \phi(\mathrm{d}x') P^{J_m} \cdots P^{J_1}(x', \mathrm{d}x).$$

For example, for the PAR scheme, the kernel $\frac{1}{2}(\mathcal{P}_{\text{odd}}\mathcal{P}_{\text{even}} + \mathcal{P}_{\text{even}}\mathcal{P}_{\text{odd}})$ is reversible. Lemma 1 can be used to define other reversible samplers.

2.3 Particle Gibbs

Our primary motivation for studying the convergence properties of blocked Gibbs samplers for general HMMs is the development of PMCMC. These methods provide systematic and efficient Markov kernels which can be used to simulate from the block-conditionals ϕ_x^J for "reasonably large" blocks, making the blocking strategy practically interesting. In particular, we will make use of the PG sampler (Andrieu et al., 2010) to define Q^J and hence the non-ideal Gibbs kernel in (4). The PG sampler of Andrieu et al. (2010) is a Markov kernel that preserves the invariance of the full JSD. In this section we briefly review the PG sampler and discuss how the standard PG algorithm needs to be modified (see Algorithm 1) when targeting one of the block-conditionals ϕ_x^J instead of the full joint smoothing distribution ϕ. (Although not detailed as in Algorithm 1, blocking was mentioned by Andrieu et al. (2010, page 294) as a possible extension of their original PG sampler. However, there was no discussion or mention of the greatly improved stability of the sampler which is our main interest here.)

The PG sampler can be used to construct a uniformly ergodic Markov kernel on $\mathsf{X}^{|J|}$ which leaves ϕ_x^J invariant. This construction is based on SMC. We denote the PG Markov kernel for

block J by Q_N^J, where N denotes a precision parameter; specifically the number of particles used in the underlying SMC sampler. We work under assumption (B-1), i.e. that J is an interval. Furthermore, for the sake of illustration, assume that J is an internal block, i.e., $J = \{s, \ldots, u\}$ for $s > 1$ and $u < n$. Obvious modifications to the algorithmic statement are needed for $s = 1$ and/or $u = n$.

The conditional probability density function of $X_{s:u}$, given $X_{s-1} = x_{s-1}$, $X_{u+1} = x_{u+1}$, and $Y_{s:t} = y_{s:t}$ is

$$\phi_x^{s:u}(x_{s:u}) = \frac{1}{p(y_{s:u} \mid x_{s-1}, x_{u+1})} \left\{ \prod_{j=s}^{u} m(x_{j-1}, x_j) g(x_j, y_j) \right\} m(x_u, x_{u+1}). \tag{6}$$

To aid the description of the PG kernel, we define for $t \in J$ the intermediate probability density

$$\pi_{x_{s-1}}^{s:t}(x_{s:t}) := p(x_{s:t} \mid x_{s-1}, y_{s:t}) = \frac{1}{p(y_{s:t} \mid x_{s-1})} \prod_{j=s}^{t} m(x_{j-1}, x_j) g(x_j, y_j), \tag{7}$$

which is the conditional probability density function of $X_{s:t}$ given $X_{s-1} = x_{s-1}$ and $Y_{s:t} = y_{s:t}$. It follows that the block-conditional density is given by

$$\phi_x^{s:u}(x_{s:u}) \propto \pi_{x_{s-1}}^{s:u}(x_{s:u}) m(x_u, x_{u+1}). \tag{8}$$

The sampling procedure used in the PG sampler is reminiscent of a standard SMC sampler, see e.g. Doucet et al. (2000); Del Moral (2004); Doucet and Johansen (2011); Cappé et al. (2005) and the references therein. Here we review a basic method, though it should be noted that the PG sampler can be generalised to more advanced methods, see Andrieu et al. (2010); Chopin and Singh (2015). As in a standard SMC sampler, the PG algorithm approximates the sequence of "target distributions" $\{\pi_{x_{s-1}}^{s:t}, t = s, \ldots, u\}$ sequentially by collections of weighted particles. The key difference between PG and standard SMC is that, in the former, one particle is set deterministically according to the current configuration of the sampler. Let this fixed particle/state trajectory (also called the reference trajectory) be denoted by $x^\star \in \mathsf{X}^n$. Then, to simulate from the PG kernel $Q_N^J(x_{J^+}^\star, \mathrm{d}x'_J)$, we proceed as follows.

Initially, we simulate particles $X_s^i \sim r_s(x_{s-1}^\star, \cdot)$, $i = 1, \ldots, N-1$ independently from some proposal density $r_s(x_{s-1}, x_s)$. The proposal density may depend on the fixed observation sequence, but we omit the explicit dependence in the notation for brevity. Note that all the particles $\{X_s^i\}_{i=1}^N$ share a common ancestor at time $s-1$, namely the fixed boundary point $x_{s-1}^\star$. Note also the we simulate only $N-1$ particles in this way. The Nth particle is set deterministically according to the current configuration: $X_s^N = x_s^\star$.

To account for the discrepancy between the proposal density and the target density $\pi_{x_{s-1}^\star}^s$, importance weights are computed in the usual way: $\omega_s^i = w_s(x_{s-1}^\star, X_s^i)$ for $i = 1, \ldots, N$, where the weight function is given by

$$w_t(x_{t-1}, x_t) := \frac{g(x_t, y_t) m(x_{t-1}, x_t)}{r_t(x_{t-1}, x_t)} \,. \tag{9}$$

The weighted particles $\{(X_s^i, \omega_s^i)\}_{i=1}^N$ provide an approximation of the target distribution $\pi_{x_{s-1}^\star}^s$, in the sense that

$$\widehat{\pi}_{x_{s-1}^\star}^s(f) := \sum_{i=1}^{N} \frac{\omega_s^i}{\sum_{\ell=1}^N \omega_s^\ell} f(X_s^i)$$

is an estimator of $\int f(x)\pi^s_{x^\star_{s-1}}(x)\mathrm{d}x$ for any measurable function $f : \mathsf{X} \to \mathbb{R}$.

We proceed inductively. Assume that we have at hand a weighted sample $\{(X^i_{s:t-1}, \omega^i_{t-1})\}_{i=1}^N$ approximating the target distribution $\pi^{s:t-1}_{x^\star_{s-1}}$ at time $t-1$. This weighted sample is then propagated sequentially *forward in time*. This is done by sampling, for each particle $i \in \{1, \ldots, N-1\}$, an *ancestor index* A^i_t with (conditional) probability

$$\mathbb{P}(A^i_t = j) = \frac{\omega^j_{t-1}}{\sum_{\ell=1}^N \omega^\ell_{t-1}}\,, \quad j \in \{1, \ldots, N\}\,. \tag{10}$$

Given the ancestor, a new particle position is sampled as $X^i_t \sim r_t(X^{A^i_t}_{t-1}, \cdot)$, where $r_t(x_{t-1}, x_t)$ is a proposal density on X. Again, note that we only generate $N-1$ particles in this way. The Nth particle and its ancestor index are set deterministically as $X^N_t = x^\star_t$ and $A^N_t = N$. The particle trajectories (i.e., the ancestral paths of the particles X^i_t, $i \in \{1, \ldots, N\}$) are constructed sequentially by associating the current particle X^i_t with the particle trajectory of its ancestor: $X^i_{s:t} := (X^{A^i_t}_{s:t-1}, X^i_t)$. It follows by construction that the Nth particle trajectory coincides with the current configuration: $X^N_{s:t} = x^\star_{s:t}$ for any $t \in J$.

Finally, the particles are assigned importance weights. For $t < u$, we compute $\omega^i_t = w_t(x^\star_{t-1}, X^i_t)$ for $i = 1, \ldots, N$, where the weight function is given by (11). At the final iteration of block J, we need to take into account the fact that the target distribution is $\phi^{s:u}_{x^\star}$ and not $\pi^{s:u}_{x^\star_{s-1}}$; the two being related according to (8). We can view the fixed boundary state $x^\star_{u+1}$ as an "extra observation" and, consequently, the weights are computed according to

$$\omega^i_u = \frac{m(X^i_u, x^\star_{u+1})g(X^i_u, y_u)m(X^{A^i_u}_{u-1}, X^i_u)}{r_u(X^{A^i_u}_{u-1}, X^i_u)} = m(X^i_u, x^\star_{u+1})w(X^{A^i_u}_{u-1}, X^i_u)\,. \tag{11}$$

Similarly to the fact that the proposal densities are allowed to depend on the fixed observations $y_{s:u}$, it is possible to let the final proposal density r_u depend on the fixed boundary point $x^\star_{u+1}$, but again we do not make this explicit in the notation for simplicity.

After a complete pass of the above procedure, the current configuration of the sampler is updated by replacing states $x^\star_J$ (recall that $J = s : u$) by a draw from the particle approximation of $\phi^J_{x^\star}$. That is, X'_J is sampled from among the particle trajectories at time u, with probabilities given by their importance weights, i.e.,

$$\mathbb{P}(X'_J = X^i_J) = \frac{\omega^i_u}{\sum_{\ell=1}^N \omega^\ell_u}\,, \quad i \in \{1, \ldots, N\}\,, \tag{12}$$

and $x^\star \leftarrow \{x^\star_{J^c}, X'_J\}$ is taken as the new configuration of the sampler. The block PG sampling procedure is summarized in Algorithm 1. Note that the procedure associates each $x^\star_{J^+}$ with a probability distribution on $\mathsf{X}^{|J|}$; this is the PG kernel Q^J_N. As shown by Andrieu et al. (2010), the conditioning on a reference trajectory implies that the PG kernel leaves the conditional distribution ϕ^J_x invariant in the sense of (3). Quite remarkably, this invariance property holds for any $N \geq 1$ (though, $N \geq 2$ is required for the kernel to be ergodic; see Andrieu et al. (2010); Lindsten et al. (2015); Andrieu et al. (2015)).

Empirically, it has been found that the mixing of the PG kernel can be improved significantly by updating the ancestor indices A^N_t, for $t \in \{1, \ldots, T\}$, either as part of a separate backward recursion (Whiteley et al., 2010), or in the the forward recursion (Lindsten et al., 2014) itself. Although we do not elaborate on the use of these modified PG algorithms in this work, the stability results of the blocked PG sampler presented in the subsequent sections also hold when the PG kernel is replaced by one of these modified algorithms (which might result in better empirical performance).

Algorithm 1 PG sampler for ϕ_x^J for $J = \{s, \ldots, u\}$ (non-boundary block)

Input: Observations y_J, fixed boundary states $x^\star_{\partial J}$, and reference states $x^\star_J$.

Output: Draw from the PG Markov kernel $Q_N^J(x^\star_{J^+}, \mathrm{d}x'_J)$.

Draw $X_s^i \sim r_s(x^\star_{s-1}, \cdot)$ for $i = 1, \ldots, N-1$ and set $X_s^N = x^\star_s$.
Set $\omega_s^i = w_s(x^\star_{s-1}, X_s^i)$ for $i = 1, \ldots, N$.
for $t = s+1$ **to** u **do**
  Draw A_t^i with $\mathbb{P}(A_t^i = j) = \omega_{t-1}^j / \sum_{\ell=1}^N \omega_{t-1}^\ell$, for $i = 1, \ldots, N-1$.
  Draw $X_t^i \sim r_t(X_{t-1}^{A_t^i}, \cdot)$ for $i = 1, \ldots, N-1$.
  Set $X_t^N = x_t^\star$ and $A_t^N = N$.
  if $t < u$ **then**
    Set $\omega_t^i = w_t(X_{t-1}^{A_t^i}, X_t^i)$ for $i = 1, \ldots, N$.
  else
    Set $\omega_u^i = m(X_u^i, x^\star_{u+1}) w_u(X_{u-1}^{A_u^i}, X_u^i)$ for $i = 1, \ldots, N$.
  end if
  Set $X_{s:t}^i = (X_{s:t-1}^{A_t^i}, X_t^i)$ for $i = 1, \ldots, N$.
end for
Draw K with $\mathbb{P}(K = j) = \omega_u^j / \sum_{\ell=1}^N \omega_u^\ell$, $j \in \{1, \ldots, N\}$.
return $X'_J := X_J^K$.

3 Main Results: Convergence of the Block Samplers

In this section we state the main convergence results for the blocked Gibbs samplers detailed in Section 2.2. After introducing our notation and stating some known preliminary results in Section 3.1, we start in Section 3.2 by deriving Wasserstein estimates for the (ideal) blocked Gibbs kernels. We then investigate in Sections 3.3 and 3.4 how these estimates result in contraction rates for the ideal and Particle Gibbs block samplers, respectively.

3.1 Preliminaries, notation, and definitions

For a function $f : \mathsf{X}^n \mapsto \mathbb{R}$, the oscillation of f with respect to the i-th coordinate is denoted by

$$\mathrm{osc}_i(f) = \sup_{\substack{x,y \in \mathsf{X}^n \\ x_{-i} = y_{-i}}} |f(x) - f(y)| \,. \tag{13}$$

while $\mathrm{osc}(f) = \sup_{x,z} |f(x) - f(z)|$ denotes the total oscillation of f. Note that

$$f(x) - f(z) \leq \sum_{i \in I} \mathrm{osc}_i(f) \mathbb{I}_{[x_i \neq z_i]}. \tag{14}$$

For a matrix A, recall that $\|A\|_\infty = \max_i \sum_j |a_{i,j}| = \max_i [A\mathbf{1}]_i$ where the last equality holds only if all elements are non-negative. The norm is sub-multiplicative, i.e. $\|AB\|_\infty \leq \|A\|_\infty \|B\|_\infty$. Let μ and ν be two probability measures on X. With Ψ being a probability measure on $\mathsf{X} \times \mathsf{X}$ we say that Ψ is a *coupling* of μ and ν if $\int \Psi(\cdot, \mathrm{d}x) = \mu(\cdot)$ and $\int \Psi(\mathrm{d}x, \cdot) = \nu(\cdot)$.

We review some well-known techniques for the analysis of Markov chains (see, e.g., Follmer (1982)). Let P be a Markov kernel on X^n. The matrix W is a *Wasserstein matrix* for P if for

any function f of finite oscillation,

$$\mathrm{osc}_j(Pf) \leq \sum_{i \in I} \mathrm{osc}_i(f) W_{ij}, \qquad \text{for all } j \in I\,. \tag{15}$$

If P and Q are two Markov kernels with Wasserstein matrices V and W, respectively, then WV is a Wasserstein matrix for the composite kernel PQ.

The convergence rate of an MCMC procedure can be characterised in terms of a corresponding Wasserstein matrix for the Markov transition kernel, through the following (well-known) result. For any probability distributions μ and ν and any function f with finite oscillation, we have

$$\begin{aligned} |\mu P^k f - \nu P^k f| &\leq \sum_{i,j \in I} \mathrm{osc}_i(f) W_{ij}^k \Psi(X_j \neq Z_j) \\ &\leq \left(\sum_{i \in I} \mathrm{osc}_i(f) \right) \|W\|_\infty^k \max_{j \in I} \Psi(X_j \neq Z_j) \end{aligned} \tag{16}$$

for any $k \geq 1$ and for any coupling Ψ of μ and ν. (Note that $\Psi(X_j \neq Z_j) = \int \Psi(\mathrm{d}x, \mathrm{d}z) \mathbb{I}_{[x_j \neq z_j]}$ is the probability of not coupling element j under the coupling Ψ.) This result can be verified by using (14) and iterating the inequality (15). It follows from (16) that $\|W\|_\infty < 1$ implies a geometric rate of contraction of the kernel P.

3.2 Wasserstein estimates for the ideal block sampler

Our convergence results, both for the ideal and for the PG samplers, rely on constructions of Wasserstein matrices for the *ideal* Gibbs kernels. Indeed, as we shall see in Section 3.4, the PG block sampler can be viewed as an ϵ-perturbation of the ideal block sampler, and this is exploited in the convergence analysis. Let P^J denote the ideal Gibbs kernel that updates block J only, as defined in (4), and let W^J be a Wasserstein matrix for P^J. The following lemma reveals the structure of W^J.

Lemma 2. *A Wasserstein matrix for the ideal Gibbs kernel updating block J can be chosen to satisfy,*

$$W_{i,j}^J = \begin{cases} \mathbb{I}_{[i=j]}, & i \in J^c, \\ 0, & i \in J,\ j \in I \setminus \partial J, \\ \times, & i \in J,\ j \in \partial J, \end{cases} \tag{17}$$

where $\times$ denotes elements which are in general in the interval $[0, 1]$.

The interpretation of components that are either 0 or 1 is noteworthy as these cannot be improved further. (Recall, the smaller the row sum of the Wasserstein matrix the better.) For example, the lemma states that when the states x and z differ in a single component, say $x_j \neq z_j$, and if $j \in J$, then this error is not propogated when computing the difference $P^J f(x) - P^J f(z)$. In this sense the given Wasserstein matrices are the 'element-wise minimal' ones.

Proof (Lemma 2). From (4) we have

$$P^J(x, \mathrm{d}x') = \phi_x^J(\mathrm{d}x'_J) \delta_{x_{J^c}}(\mathrm{d}x'_{J^c})\,.$$

A candidate Wasserstein matrix W^J can be found via a coupling argument. For any pair (x, z) such that $x_{-j} = z_{-j}$ and $x_j \neq z_j$, let $\Psi^J_{j,x,z}$ be a coupling of ϕ^J_x and ϕ^J_z. Consider now

$$\begin{aligned}
|P^J f(x) - P^J f(z)| &\leq \int \Psi^J_{j,x,z}(\mathrm{d}x'_J, \mathrm{d}z'_J)\delta_{x_{J^c}}(\mathrm{d}x'_{J^c})\delta_{z_{J^c}}(\mathrm{d}z'_{J^c})|f(x') - f(z')| \\
&\leq \sum_{i\in I} \operatorname{osc}_i(f) \int \Psi^J_{j,x,z}(\mathrm{d}x'_J, \mathrm{d}z'_J)\delta_{x_{J^c}}(\mathrm{d}x'_{J^c})\delta_{z_{J^c}}(\mathrm{d}z'_{J^c})\mathbb{I}_{[x'_i \neq z'_i]} \\
&= \sum_{i\in J} \operatorname{osc}_i(f) \int \Psi^J_{j,x,z}(\mathrm{d}x'_i, \mathrm{d}z'_i)\mathbb{I}_{[x'_i \neq z'_i]} + \sum_{i\in J^c} \operatorname{osc}_i(f)\mathbb{I}_{[x_i \neq z_i]} \\
&= \sum_{i\in J} \operatorname{osc}_i(f)\Psi^J_{j,x,z}(X'_i \neq Z'_i) + \mathbb{I}_{[j\in J^c]} \operatorname{osc}_j(f) \ ,
\end{aligned}$$

where the second line follows from (14) and the last line follows since $x_{-j} = z_{-j}$, i.e., at most one term from the second sum in the penultimate line will be non-zero.

Thus, for $i \in J$ we can set

$$W^J_{i,j} := \sup_{\substack{x,\, z \,\in\, \mathsf{X}^n \\ x_{-j} \,=\, z_{-j}}} \Psi^J_{j,x,z}(X'_i \neq Z'_i). \tag{18}$$

and for $i \notin J$, $W^J_{i,j} := \mathbb{I}_{[i=j\in J^c]}$. Furthermore, since ϕ^J_x depends on x only through the boundary points $x_{\partial J}$, it follows that for $j \notin \partial J$ the coupling Ψ_j can be made perfect: $\Psi_{j,x,z}(X'_i = Z'_i) = 1$ for any $i \in J$. Therefore, it is evident from (18) that $W^J_{i,j} = 0$ for $i \in J$ and $j \in I \setminus \partial J$. □

Specifically, if $J = s : u$ is an interval (cf.(B-1)) it follows that W^J is structured as (blanks correspond to zeros),

$$W^J = \begin{bmatrix}
1 & & & & & & & & \\
& \ddots & & & & & & & \\
& & 1 & & & & & & \\
& & \times & 0 & \cdots & 0 & \times & & \\
& & \vdots & \vdots & \ddots & \vdots & \vdots & & \\
& & \times & 0 & \cdots & 0 & \times & & \\
& & & & & & 1 & & \\
& & & & & & & \ddots & \\
& & & & & & & & 1
\end{bmatrix}, \tag{19}$$

with obvious modifications for the "boundary blocks" ($s = 1$ and/or $u = n$). The "square of zeros" correspond to rows and columns $i \in J$ and $j \in J$, respectively. The columns of $\times$'s correspond to $W^J_{i,j}$ with $i \in J$ and $j \in \partial J$.

It remains to compute the non-zero, off-diagonal elements of W^J. To this end, we use the strong, but standard, mixing conditions (S-1) and (S-2), given in Section 2.1. This allows us to express the "unknown" elements of W^J in terms of the mixing coefficients of the model.

Lemma 3. *Assume* (B-1) *and let* $J = s : u$. *Assume* (S-1) *and* (S-2) *and define the constant* $\alpha \in [0, 1)$ *as*

$$\alpha := 1 - \delta^{\frac{1-h}{h}} \frac{\sigma_-}{\sigma_+} \ , \tag{20}$$

where δ, σ_-, σ_+, *and* h *are defined in* (S-1) *and* (S-2). *Then,* W^J *defined as in* (17) *and with*

$$W^J_{i,j} = \begin{cases} \alpha^{\lfloor h^{-1}(i-j)\rfloor}, & j = s-1 \ \textit{(if } s > 1\textit{)}, \\ \alpha^{\lfloor h^{-1}(j-i)\rfloor}, & j = u+1 \ \textit{(if } u < n\textit{)}, \end{cases}$$

for $i \in J$ and $j \in \partial J$ is a Wasserstein matrix for the ideal Gibbs block-transition kernel P^J.

Proof. See Appendix B. □

3.3 Contraction of the ideal block sampler

Let $\partial := \bigcup_{J \in \mathcal{J}} \partial J$ denote the set of all boundary points. We start by stating a general geometric convergence result which holds for an arbitrary cover $\mathcal{J}$ of I.

(A-1) For all $J \in \mathcal{J}$, $\max_{i \in J \cap \partial} \sum_{j \in \partial J} W^J_{i,j} \leq \lambda < 1$.

Theorem 4. *Assume* (A-1) *and let W^J satisfy* (17) *for all $J \in \mathcal{J}$. Let $\mathcal{J} = \{J_1, \ldots, J_m\}$ be an* arbitrary cover *of I and let $\mathcal{W} := W^{J_m} \cdots W^{J_1}$ be the Wasserstein matrix for $\mathcal{P} = P^{J_1} \cdots P^{J_m}$, i.e., for one complete sweep of the ideal Gibbs sampler. Then, for $k \geq 1$, $\|\mathcal{W}^k\|_\infty \leq \lambda^{k-1} \|\mathcal{W}\|_\infty$.*

Proof. See Appendix A. □

The following result now characterises the convergence of the law of the sampled output of the ideal block sampler.

Corollary 5. *Let μ and ν be two probability distributions on X^n and let Ψ be an arbitrary coupling of μ and ν. Under the same conditions as in Theorem 4 we have, for any $k \geq 1$ and any f of finite oscillation,*

$$|\mu \mathcal{P}^k f - \nu \mathcal{P}^k f| \leq \|\mathcal{W}\|_\infty \lambda^{k-1} \max_{j \in I} \{\Psi(X_j \neq Z_j)\} \sum_{i \in I} \mathrm{osc}_i(f).$$

Corollary 5 verifies two important facts. First is that if we are interested only in the convergence of certain marginals of the sampled process $\{X[k], k \in \mathbb{N}\}$ in (2) to the corresponding marginal of the JSD, then the convergence rate is *independent* of the dimension n of the JSD. Secondly, for convergence in total variation norm of the law of $X[k]$ to the full JSD, we attain a bound on the error which is $O(n\lambda^k)$. That is, the error grows slowly (linearly) with the dimension n of the JSD.

Under the conditions of Lemma 3 we can clearly see the benefit of blocking for verifying condition (A-1). As an illustration, let $h = 1$ in (S-1). The condition for contraction requires that for any $i \in J \cap \partial$ (assuming J is an internal block),

$$\sum_{j \in \partial J} W^J_{i,j} = \alpha^{i-(s-1)} + \alpha^{(u+1)-i} < 1 \,, \tag{21}$$

where $\alpha \in [0, 1)$ is defined in Lemma 3. First of all, we note that it is possible to ensure $\sum_{j \in \partial J} W^J_{i,j} < 1$ for *any* $i \in J$ by increasing the block size $L := u - s + 1$. Indeed, the maximum of (21) for $i \in J$ is attained for $i = s$ or $i = u$, for which $\sum_{j \in \partial J} W^J_{i,j} = \alpha + \alpha^L$. Secondly, however, we note that Lemma 3 also reveals the benefit of using *overlapping* blocks. Indeed, since we only need to control (21) for $i \in J \cap \partial$, we can select the blocking scheme so that the set of boundary points ∂ excludes indices i close to the boundary of block J. Consequently, by using overlapping blocks we can control both terms in (21)—and thus the overall convergence rate of the algorithm—by increasing L.

Theorem 4 assumes no specific structure for $\mathcal{J}$ other than it being a cover. As such, it cannot provide a sharper estimate of the contraction since it caters for all blocking structures. In order to refine the contraction estimate we impose the blocking structure formalised by (B-1) and (B-2), as illustrated in Figure 1, and study the interplay between block size, overlap, and convergence. The theorem below improves the estimate of the decay of errors per complete sweep from λ in Theorem 4 to λ^2 for the blocking structure of Figure 1.

Theorem 6. *Assume* (B-1), (B-2), *and* (A-1). *Then, for any* $k \geq 1$*:*

- *For the ideal PAR sampler,*
$$\|\mathcal{W}^k\|_\infty \leq \|\mathcal{W}\|_\infty \lambda^{2(k-1)} \tag{22}$$
and $\|\mathcal{W}\|_\infty \leq 2$*, where* λ *is defined as in* (A-1).

- *For the ideal L-R sampler,*
$$\|\mathcal{W}^k\|_\infty \leq \|\mathcal{W}\|_\infty \beta^{k-1} \tag{23}$$
and $\|\mathcal{W}\|_\infty \leq 1+\lambda$*, where* $\beta = \max_{k\in 2:m} \lambda a_k + b_k$*,* $a_k = W^{J_k}_{\partial_+ J_{k-1}, \partial_- J_k}$*,* $b_k = W^{J_k}_{\partial_+ J_{k-1}, \partial_+ J_k}$.

Proof. See Appendix A. □

Remark 7. *There is parity in the two rates of Theorem 6 since it can be shown that* $\beta \approx \lambda^2$*. For example let each block be the same length* L *and the overlap between all adjacent blocks* $J_{k-1}, J_k \in \mathcal{J}$ *be fixed,* $|J_{k-1} \cap J_k| = p$*. Under strong mixing* (S-1)–(S-2)*,* $b_k = \alpha^{\lfloor h^{-1}(L-p)\rfloor}$*,* $a_k = \alpha^{\lfloor h^{-1}(p+1)\rfloor}$*,* $\lambda = \alpha^{\lfloor h^{-1}(L-p)\rfloor} + \alpha^{\lfloor h^{-1}(p+1)\rfloor}$ *and* $\beta/\lambda^2 \to 1$ *as* L *increases and* p *is fixed.*

3.4 Contraction of the Particle Gibbs block sampler

We now turn our attention to the Particle Gibbs block sampler. In this section we state a main result (Theorem 8) that parallels Theorem 6 for the PG kernel. For the sake of interpretability, we specialize the result to the case of a common block size L and common overlap p between successive blocks (see also Remark 7). A version of Theorem 8 without this assumption, nor strong mixing, is presented in Section 4; see Theorems 9 and 10. (Theorem 8 is a corollary of these theorems.)

(B-3) For all $J \in \mathcal{J}$, $|J| = L$ and for all consecutive $J_{k-1}, J_k \in \mathcal{J}$, $|J_{k-1} \cap J_k| = p$.

Note that (B-3) implies that n is assumed to satisfy $n = (L-p)m + p$.

Theorem 8. *Assume* (B-1)–(B-3) *and* (S-1)–(S-2). *Let* $\mathcal{P}$ *denote the Markov kernel corresponding to one complete sweep of either the PAR sampler or the L-R sampler, and assume that each block is updated by simulating from the PG kernel (Algorithm 1) using a bootstrap proposal:* $r(x, x') = m(x, x')$*. Let* μ *and* ν *be two probability distributions on* X^n *and let* Ψ *be an arbitrary coupling of* μ *and* ν*. Then, for any* f *of finite oscillation and any* $k \geq 1$

$$|\mu\mathcal{P}^k f - \nu\mathcal{P}^k f| \leq \lambda^k_{\mathrm{PG}} \times \Psi(X \neq Z) \sum_{i\in I} \mathrm{osc}_i(f), \tag{24}$$

where:

- *For the Particle Gibbs PAR sampler,*
$$\lambda_{\mathrm{PG}} \leq \lambda(\beta \vee 1) + \epsilon\left(2\lambda + 25\epsilon + 8(\beta \vee 1)\right). \tag{25}$$

- *For the Particle Gibbs L-R sampler,*
$$\lambda_{\mathrm{PG}} \leq \lambda + \alpha^{\lfloor h^{-1}(L-p+1)\rfloor} + 2\epsilon \frac{3(\beta \vee 1) + 1 + \lambda}{1 - 2\epsilon - \alpha^{\lfloor h^{-1}(p+1)\rfloor}} \tag{26}$$
provided $2\epsilon + \alpha^{\lfloor h^{-1}(p+1)\rfloor} < 1$*,*

provided that $\lambda < 1$. *In the above,*

$$\lambda = 2\alpha^{\lfloor h^{-1}(p+1)\rfloor}, \qquad \beta = \alpha^{\lfloor h^{-1}\rfloor} + \alpha^{\lfloor h^{-1}(L-p+1)\rfloor}, \tag{27}$$

and $\epsilon = \epsilon(N, L) = 1 - (1 - \frac{1}{c(N-1)+1})^L$, *for some constant* c *(specified in Proposition 12 below) which is independent of* n, N, L *and* p.

Proof. The proof is given in Section 4. □

Theorem 8 also applies to the ideal sampler (set $\epsilon = 0$). In terms of sufficiency for contraction, the requirement that the non-ϵ terms of this theorem be less than 1 is stronger than (A-1); this should not be surprising since the analysis is catered for the non-ideal PG kernel and thus is inherently more conservative. For common blocks lengths L and overlap p, (A-1) requires $\alpha^{\lfloor \frac{p+1}{h}\rfloor} + \alpha^{\lfloor \frac{L-p}{h}\rfloor} < 1$. The non-$\epsilon$ terms of Theorem 8 are shrunk by increasing the overlap of blocks and then increasing block size with overlap fixed. Alternatively, if p is a constant fraction of L, then λ tends to zero as L increases. We remark that the non-ϵ rates of Theorem 8 are not as sharp as they could be as given in Theorem 9 and Theorem 10 below. The upper bounds for the non-ϵ terms were chosen for the sake of simplicity and interpretability since they are functions of system forgetting (α and h), common block lengths L, and overlap p only.

4 Proof of the Contraction of the Particle Gibbs Block Sampler

This section is dedicated to the proof of Theorem 8. Furthermore, it provides a more general version of Theorem 8 that avoids the common block length and overlap structure of assumption (B-3). The general results given below are stated in terms of Wasserstein estimates for the *non-ideal* blocked Gibbs sampler $\widehat{W}$, under the assumption (A-2) that $\widehat{W}$ is an ϵ-perturbation of a Wasserstein matrix for the ideal block sampler. Specifically, let $\widehat{W}^J$ be a Wasserstein matrix for the non-ideal block-transition kernel defined in (5). By an analogous argument as in Lemma 2 it follows that $\widehat{W}^J$ has a similar structure as W^J, but with (possibly) non-zero entries also for rows and columns $i \in J$ and $j \in J$, respectively, which motivates the following assumed structure.

(A-2) There exists a (common) constant $\epsilon \in [0, 1)$ and, for each $J \in \mathcal{J}$, a matrix W^J satisfying (17) such that, for any $J \in \mathcal{J}$, $\widehat{W}^J_{i,j} := W^J_{i,j} + \epsilon\, \mathbb{I}_{[i\in J, j\in J_+]}$ is a Wasserstein matrix for the non-ideal block-transition kernel updating block J.

In Proposition 12 below we show that (A-2) indeed holds for the PG kernel, with W^J being a Wasserstein matrix for block J of the ideal sampler (as in Lemma 3) and where the perturbation ϵ depends on (grows with) block size $|J|$ and (decreases with) the number of particles N, but is independent of the length of the total data record n or the specific HMM observations pertaining to each block. This is key to the stability of the PG block sampler as $n \to \infty$ for fixed N.

Theorem 9. *Assume* (B-1), (B-2), (A-2), *PAR and assume that the number of blocks* $m = |\mathcal{J}|$ *is odd.*[1] *Let* $J_{-k} = \cup_{J\in\{\mathcal{J}\setminus J_k\}} J$ *and* $L = \max_{J\in\mathcal{J}} |J|$. *Then* $\widehat{\mathcal{W}}$, *the Wasserstein matrix of one complete sweep (defined analogously to Theorem 4,) satisfies*

$$[\widehat{\mathcal{W}}\mathbf{1}]_i \le \begin{cases} \lambda^2 + \epsilon\left(\lambda(L+4) + \epsilon\,(L+2)^2 + L(1\vee\beta)\right), & i \in J^c_{-k},\ k \text{ even}, \\ \lambda + \epsilon\,(L+2), & i \in J^c_{-k},\ k \text{ odd}, \\ \lambda\beta + \epsilon\left(\beta(L+2) + 2\lambda + \epsilon\,(L+2)^2 + L(1\vee\beta)\right), & i \in J_k \cap J_{-k},\ k \text{ even}, \end{cases}$$

[1] Clearly, an analogous result holds for even m.

where

$$\begin{aligned}\lambda &= \max_{J_k\in\mathcal{J}} \max_{i\in J^c_{-k}} W^{J_k}_{i,\partial_- J_k} + W^{J_k}_{i,\partial_+ J_k},\\ \beta &= \max_{J\in\mathcal{J}} \max_{i\in J} \sum_{j\in\partial J} W^J_{i,j}\end{aligned}$$

and $W^{J_1}_{i,\partial_- J_1} = W^{J_m}_{i,\partial_+ J_m} = 0$ *by convention.*

Proof. See Appendix A. □

Theorem 10. *Assume* (B-1), (B-2), (A-2) *and L-R. Let* $J_{-k} = \cup_{J\in\{\mathcal{J}\setminus J_k\}} J$, $L = \max_{J\in\mathcal{J}} |J|$ *and* $L_1 = \max_{k\in 2:m} |J_{k-1}\cap J_k|$. *Then* $\widehat{\mathcal{W}}$ *satisfies*

$$[\widehat{\mathcal{W}}\mathbf{1}]_i \le \begin{cases}\lambda + c\epsilon, & i\in J^c_{-k},\ k\in 1:m,\\ \lambda' + 2c\epsilon, & i\in J_{k-1}\cap J_k,\ k\in 2:m,\end{cases}$$

where,

$$\begin{aligned}\lambda &= \max_{J_k\in\mathcal{J}} \max_{i\in J^c_{-k}} \frac{W^{J_k}_{i,\partial_+ J_k}}{1 - W^{J_k}_{i,\partial_- J_k}},\\ \lambda' &= \max_{k\in 2:m} \max_{i\in J_{k-1}\cap J_k} \lambda W^{J_k}_{i,\partial_- J_k} + W^{J_k}_{i,\partial_+ J_k},\\ c &= \frac{L(\beta(\lambda\vee 1)\vee 1) + 1 + \lambda}{1 - (L_1+1)\epsilon - \gamma},\\ \beta &= \max_{J\in\mathcal{J}} \max_{i\in J} \sum_{j\in\partial J} W^J_{i,j},\\ \gamma &= \max_{k\in 2:m} \max_{i\in J_k\cap J^c_{k-1}} W^{J_k}_{i,\partial_- J_k},\end{aligned}$$

provided that $\gamma + (L_1+1)\epsilon < 1$. *(By convention set* $W^{J_1}_{i,\partial_- J_1} = W^{J_m}_{i,\partial_+ J_m} = 0$.*)*

Proof. See Appendix A. □

In order to prove Theorem 8 we can now make use of Lemma 3 (assuming (S-1)–(S-2) and (B-1)–(B-3)) to identify the constants of Theorems 9 and 10. However, a technical detail is to handle the dependence on the maximum block size $L = \max_{J\in\mathcal{J}} |J|$ and maximum overlap $L_1 = \max_{k\in 2:m} |J_{k-1}\cap J_k|$ (of Theorem 10). Indeed, a direct application of Theorems 9 and 10 would suggest that the norm of $\widehat{\mathcal{W}}$ grows, respectively, quadratically or linearly with ϵL. To avoid this issue we will make use of the following trick: when applying Theorems 9 and 10 we do not consider the original HMM formulation, but an equivalent model that lumps consecutive states together (for the sake of the analysis), thus effectively reducing the size of the blocks. Under (B-1)–(B-3), each block can be split into three distinct sections which are the left overlap, the middle of the block, and the right overlap. (See Figure 1.) The exception are the end blocks which are split into two sections. By viewing each of these parts as a single lumped state, we reduce the block size to 3 and maximum overlap to 1. For this scheme, the lumped states are given by

$$\begin{aligned}&\Xi_1 = X_{1:L-p}, && \Xi_2 = X_{L-p+1:L},\\ &\Xi_3 = X_{L+1:2L-2p}, && \Xi_4 = X_{2L-2p+1:2L-p}, && \Xi_5 = X_{2L-p+1:3L-3p},\\ &\ \vdots\end{aligned}$$

where Ξ_1 and Ξ_2 are the states of the two sections of block 1, Ξ_2, Ξ_3 and Ξ_4 are the states of block 2 etc. More precisely:

Remark 11. *The Ξ-system groups the random variables $X_1, \ldots, X_n$ of the X-system as $\Xi_1, \ldots, \Xi_{2m-1}$, noting that* (B-3) *implies that $n = (L-p)m + p$, where*

$$\begin{cases} \Xi_1 = X_{1:L-p}, & \\ \Xi_{2i} = X_{(L-p)i+1:(L-p)i+p}, & 1 \le i < m, \\ \Xi_{2i-1} = X_{(L-p)(i-1)+p+1:(L-p)i}, & 1 < i < m, \\ \Xi_{2m-1} = X_{(L-p)(m-1)+p+1:(L-p)m+p}. & \end{cases}$$

The index set for the Ξ-system is $I_\Xi = \{1, \ldots, 2m-1\}$ and the cover $\mathcal{J}_\Xi$ of I_Ξ has m sets with (set k) $J_{\Xi,k} = \{2k-2, 2k-1, 2k\} \cap I_\Xi$.

To find a Wasserstein matrix for the Ξ-system we note that any conditional density of the states Ξ_i, $i \in J_{\Xi,k}$, conditionally on the boundaries of the block $J_{\Xi,k}$ and the observation pertaining to that block, is coupled analogously to the X-system; see the proof of Lemma 3 in the appendix. Analogously to Lemma 3, a Wasserstein matrix for block $J_{\Xi,k}$ of the Ξ-system is thus given by the $(2m-1) \times (2m-1)$ matrix

$$W_\Xi^k = \begin{bmatrix} 1 & & & & & & & & & \\ & \ddots & & & & & & & & \\ & & 1 & & & & & & & \\ & & \alpha^{\lfloor h^{-1} \rfloor} & 0 & 0 & 0 & \alpha^{\lfloor h^{-1}(L-p+1) \rfloor} & & & \\ & & \alpha^{\lfloor h^{-1}(p+1) \rfloor} & 0 & 0 & 0 & \alpha^{\lfloor h^{-1}(p+1) \rfloor} & & & \\ & & \alpha^{\lfloor h^{-1}(L-p+1) \rfloor} & 0 & 0 & 0 & \alpha^{\lfloor h^{-1} \rfloor} & & & \\ & & & & & & 1 & & & \\ & & & & & & & \ddots & & \\ & & & & & & & & 1 & \end{bmatrix}, \tag{28}$$

where the "square of zeros" correspond to rows/columns $2k-2$, $2k-1$, and $2k$ (with obvious modifications for $k = 1$ or $k = 2m-1$). The proof is omitted for brevity, but follows analogously to the proof of Lemma 3.

As a final ingredient to prove Theorem 8 we need to verify condition (A-2) for the PG kernel.

Proposition 12. *Assume* (S-1), (S-2) *and assume that the bootstrap proposal kernel is used in Algorithm 1: $r(x, x') = m(x, x')$. Then,*

1. *For any $J = \{s, \ldots, u\} \subseteq I$ with $u \ge s$ and any $N \ge 2$,*

$$Q_N^J(x_{J_+}, \mathrm{d}x'_J) \ge (1 - \epsilon(N, |J|))\, \phi_x^J(\mathrm{d}x'_J)$$

where $\epsilon(N, L) = 1 - (1 - \frac{1}{c(N-1)+1})^L$,

$$c = \left(2\delta \frac{\sigma_+}{\sigma_-} - 1 \right)^{-1} \in (0, 1]$$

and δ, σ_- and σ_+ are defined in (S-1), (S-2).

2. *Assume further* (B-1)–(B-3). *For $k = 1, \ldots, m$, let Let W_Ξ^k be given by* (28) *(i.e., W_Ξ^k is a Wasserstein matrix, based on the Ξ-system, for the ideal Gibbs kernel updating block J_k). Then $\widehat{W}_\Xi^k = [\widehat{W}_{\Xi,i,j}^k]_{i,j \in 1:2m-1}$ with*

$$\widehat{W}_{\Xi,i,j}^k = W_{\Xi,i,j}^k + \epsilon(N, |J_k|) \mathbb{I}_{\left[\begin{smallmatrix} 2k-2 \le i \le 2k \\ 2k-3 \le j \le 2k+1 \end{smallmatrix} \right]} \tag{29}$$

is a Wasserstein matrix for the PG kernel $Q_N^{J_k}$.

Proof. See Appendix B. □

Remark 13. *We emphasise that the lumping of state variables is used only for the sake of analysis, to improve the contraction rates in Theorems 9 and 10 by avoiding a dependence on the block size. For all practical purposes the lumping has no effect: when implementing the PG kernel (Algorithm 1) we still use the original state variables. Consequently, lumping does not affect the ergodicity or the convergence rate of the PG kernel. In particular, note that* ϵ *in* (29) *depends on the size of block* k *as expressed in the* X*-system (*$|J_k|$*) and not in the lumped* Ξ*-system (*$|J_{\Xi,k}|$*).*

We conclude this section by putting the pieces together to prove Theorem 8.

Proof (Theorem 8). Theorem 8 is established by applying Theorems 9 and 10 to the Ξ-system. The bound (24) is established by using (16) for the Ξ-system:

$$|\mu\mathcal{P}^k f - \nu\mathcal{P}^k f| \leq \left(\sum_{i=1}^{2m-1} \sup_{\{\Xi,\check{\Xi}\in\mathsf{X}^n:\Xi_{-i}=\check{\Xi}_{-i}\}} |f(\Xi) - f(\check{\Xi})|\right) \|\mathcal{W}_\Xi\|_\infty^k \max_{j\in 1:2m-1} \Psi(\Xi_j \neq \check{\Xi}_j).$$

Since $(\Xi_1, \ldots, \Xi_{2m-1}) \in \mathsf{X}^n$, we can crudely bound the sum with $\sum_{i=1}^n \sup_{\{x,y\in\mathsf{X}^n:x_{-i}=y_{-i}\}} |f(x)-f(y)|$. The final factor is also crudely bounded by $\max_{j\in I} \Psi(\Xi_j \neq \check{\Xi}_j) \leq \Psi(\Xi \neq \check{\Xi})$, for any coupling Ψ of μ and ν with $(\Xi, \check{\Xi}) \sim \Psi$.

Now it remains to derive the expression for λ_{PG} as an upper bound on $\|\mathcal{W}_\Xi\|_\infty$ for both cases PAR and L-R. We detail the case PAR only, as L-R follows by analogous arguments. For case PAR we use Theorem 9 for the Ξ-system. Theorem 9 is applicable for the Ξ-system since the Ξ-system satisfies (B-1), (B-2) (see Remark 11) and (A-2). The fact that the Ξ-system satisfies (A-2) follows from Proposition 12. Each block $J_{\Xi,k}$ of the Ξ-system has 3 elements (except the initial and final blocks which both have 2 elements each); see Remark 11. The specific values of λ and β of Theorem 9 follow from this simple 3-element block structure and the declared Wasserstein matrix in (28). The coefficient of the ϵ-term in (25) follows from a trivial bound of the three separate ϵ-coefficients given in Theorem 9 using $L = 3$. □

A Proofs of the Main Theorems

Lemma 14. *Let* $\mathcal{J} = \{J_1, \ldots, J_m\}$ *be an arbitrary cover of* I*. For each* $J \in \mathcal{J}$*, let* W^J *be a matrix with structure* (17) *and let* $\mathcal{W} = W^{J_m} \cdots W^{J_1}$*. For* $j \in I$*, define,* $a_j := \min\{k : j \in J_k\}$ *and* $b_j = \min\{k : j \in \partial J_k\}$*, with the convention that* $\min\{\emptyset\} = \infty$*. If* $a_j < b_j$*, then* $\mathcal{W}e_j = 0$*.*

Proof. Due to the structure of W^J from (17) it follows that, for $j \notin \partial J$, $W^J e_j = \mathbb{I}_{[j\in J^c]} e_j$. If $a_j < b_j$, then

$$\begin{aligned}\mathcal{W}e_j &= W^{J_m} \cdots W^{J_{b_j}} \cdots W^{J_{a_j}} \cdots W^{J_1} e_j \\ &= W^{J_m} \cdots W^{J_{b_j}} \cdots W^{J_{a_j}} e_j = 0,\end{aligned}$$

where the first equality follows from the fact that for any $k < a_j < b_j$, $j \notin J_k \cup \partial J_k$ by the definition of a_j and b_j. □

A.1 Proof of Theorem 4

From Lemma 14 it is clear that $\mathcal{W}e_j = 0$ for any $j \notin \partial$ (for which $b_j = \infty$). Hence, with M being the binary mask matrix with elements $M_{i,j} = \mathbb{I}_{[i=j\in\partial]}$, it holds that $\mathcal{W}r = \mathcal{W}(Mr)$ for any vector r. Thus $\mathcal{W}\mathcal{W}^{k-1}r = \mathcal{W}(M\mathcal{W})^{k-1}r$ and hence $\|\mathcal{W}^k\|_\infty \le \|\mathcal{W}\|_\infty \|M\mathcal{W}\|_\infty^{k-1}$.

Next, we consider $\|M\mathcal{W}\|_\infty = \max_{i\in\partial}\left\{e_i^\mathsf{T}\mathcal{W}\mathbf{1}\right\}$. Define recursively, $L_0 := \emptyset$ and $r_0 := \mathbf{1}$ and

$$\begin{cases} L_k := L_{k-1} \cup J_k, \\ r_k := W^{J_k} r_{k-1}, \end{cases} \quad \text{for } k = 1, \ldots, m. \tag{30}$$

We thus have $r_m = \mathcal{W}\mathbf{1}$. Hence, the result follows if, for $k = 0, \ldots, m$,

$$[r_k]_i \le \begin{cases} 1 & \text{if } i \in I \setminus L_k, \\ \lambda & \text{if } i \in L_k \cap \partial, \end{cases} \tag{31}$$

(note that nothing is said about $[r_k]_i$ for $i \in L_k \setminus \partial$). Indeed, the fact that $\mathcal{J}$ is a cover of I implies that $L_m = I$. Hence $[\mathcal{W}\mathbf{1}]_i = [r_m]_i \le \lambda$ for all $i \in \partial$.

It remains to prove (31). For $k = 0$ the hypothesis is true by construction. We proceed inductively. Hence, assume that the hypothesis is true for $k-1$. Consider

$$[r_k]_i = [W^{J_k} r_{k-1}]_i = \begin{cases} [r_{k-1}]_i & \text{if } i \in J_k^c, \\ \sum_{j\in\partial J_k} W^{J_k}_{i,j} [r_{k-1}]_j & \text{if } i \in J_k, \end{cases} \tag{32}$$

where we have made use of the structure of the Wasserstein matrix from Lemma 2. We need to consider different cases:

1. $i \in I \setminus L_k \Rightarrow [r_k]_i = [r_{k-1}]_i \le 1$, where we first use (32) and then the induction hypothesis, and the fact that $L_k = L_{k-1} \cup J_k$.
2. $i \in (L_{k-1} \setminus J_k) \cap \partial \Rightarrow [r_k]_i = [r_{k-1}]_i \le \lambda$, where, again, we use (32) and the induction hypothesis.
3. $i \in J_k \cap \partial \Rightarrow [r_k]_i = \sum_{j\in\partial J} W^{J_k}_{i,j} [r_{k-1}]_j \le \sum_{j\in\partial J} W^{J_k}_{i,j} \le \lambda$, where we use the fact that for $j \in \partial$, $[r_{k-1}]_j \le 1$ and assumption (A-1) for the final inequality.

This completes the proof. □

A.2 Proof of Theorem 6

Case L-R: The proof is similar to that of Theorem 4, but to exploit the structure of the L-R sampler we define the mask matrix M as $M_{i,j} = \mathbb{I}_{[i=j\in\partial_+]}$ where $\partial_+ = \cup_{J\in\mathcal{J}} \partial_+ J$ is the set of *right boundary points*, only, of all blocks. If $j = \partial_- J_k$ is the left boundary of some block k, say, then $j \in J_{k-1}$ by definition of the L-R sampler. Hence, from Lemma 14 it follows that $\mathcal{W}r = \mathcal{W}(Mr)$ for any vector r. Thus, $\|\mathcal{W}^k\|_\infty \le \|\mathcal{W}\|_\infty \|M\mathcal{W}\|_\infty^{k-1}$.

Define L_k and r_k as in (30). Note that (31) and (32) hold for any cover $\mathcal{J}$, and in particular for the L-R sampler. Thus, for $i \in J_k$,

$$[r_k]_i = W^{J_k}_{i,\partial_- J_k} [r_{k-1}]_{\partial_- J_k} + W^{J_k}_{i,\partial_+ J_k} [r_{k-1}]_{\partial_+ J_k} \le W^{J_k}_{i,\partial_- J_k} \lambda + W^{J_k}_{i,\partial_+ J_k}\,, \tag{33}$$

where the inequality follows from the fact that, for the L-R sampler, $\partial_- J_k \in L_{k-1}$ and $\partial_+ J_k \in I \setminus L_k$ and by using (31). Hence, $[r_k]_{\partial_+ J_{k-1}} \le \beta$. Since, for any $k \in 1:m$, $\partial_+ J_{k-1}$ lies in at

most one block (namely J_k), we conclude that $[r_m]_i \leq \beta$ for any $i \in \partial_+$. Thus $\|M\mathcal{W}\|_\infty \leq \beta$. The bound $\|\mathcal{W}\|_\infty \leq 1 + \lambda$ follows similarly from (33).

Case PAR: For the PAR sampler we redefine the mask M to be $M_{i,j} = \mathbb{I}_{[i=j\in\partial_{odd}]}$, where ∂_{odd} is the set of all boundary points of odd blocks, i.e., $\cup_{k \text{ odd}} \partial J_k$. Since any boundary point of an even block is the interior of some odd block, it follows from Lemma 14 that $\mathcal{W}r = \mathcal{W}(Mr)$ for any vector r.[2] To complete the proof we thus need to bound $\|M\mathcal{W}\|_\infty \leq \lambda^2$. Let $i \in \partial_{\text{odd}}$ and let J_k be the block such that $i \in J_k$. Note that k is even and J_k is the only block containing i. Thus, $[\mathcal{W}\mathbf{1}]_i = [\mathcal{W}_{even}\mathcal{W}_{odd}\mathbf{1}]_i = [W^{J_k}\mathcal{W}_{odd}\mathbf{1}]_i$. Now assume $[\mathcal{W}_{odd}\mathbf{1}]_i \leq \lambda$ for all i which are boundaries of even blocks (a fact we will prove next). Thus

$$[W^{J_k}\mathcal{W}_{odd}\mathbf{1}]_i = W^{J_k}_{i,\partial_- J_k}[\mathcal{W}_{odd}\mathbf{1}]_{\partial_- J_k} + W^{J_k}_{i,\partial_+ J_k}[W^{J_k}\mathcal{W}_{odd}\mathbf{1}]_{\partial_+ J_k} \leq \lambda^2.$$

To conclude: let i be a boundary of an even block and let $i \in J_k$, the only odd block containing i. Then $[\mathcal{W}_{odd}\mathbf{1}]_i = [W^{J_k}\mathbf{1}]_i = W^{J_k}_{i,\partial_- J_k} + W^{J_k}_{i,\partial_+ J_k} \leq \lambda$. □

A.3 Proof of Theorem 9

We prove the following more general version (while Theorem 9 was stated for $r = \mathbf{1}$).

Lemma 15. *Let vector r satisfy*

$$[r]_i = \begin{cases} a, & i \in J^c_{-k},\ k \text{ even}, \\ a', & i \in J^c_{-k},\ k \text{ odd}, \\ b, & i \in J_k \cap J_{-k},\ k \text{ even}, \end{cases}$$

for some positive constants a, a', b. Then

$$[\widehat{\mathcal{W}}r]_i \leq \begin{cases} \lambda^2 a + \epsilon\left(\lambda c + 2\lambda a + 2\epsilon c + La(1 \vee \beta) + \epsilon Lc\right), & i \in J^c_{-k},\ k \text{ even}, \\ \lambda a + \epsilon\left(2a + L(a' \vee b)\right), & i \in J^c_{-k},\ k \text{ odd}, \\ \lambda\beta a + \epsilon\left(\beta c + 2\lambda a + 2\epsilon c + La(1 \vee \beta) + \epsilon Lc\right), & i \in J_k \cap J_{-k},\ k \text{ even}, \end{cases}$$

where $c = 2a + L(a' \vee b)$ and the remaining constants were defined in Theorem 9.

Proof. Note that $\widehat{\mathcal{W}} = \widehat{\mathcal{W}}_{\text{even}}\widehat{\mathcal{W}}_{\text{odd}}$ where

$$\widehat{\mathcal{W}}_{\text{odd}} = \widehat{W}^{J_m} \cdots \widehat{W}^{J_3}\widehat{W}^{J_1} \qquad \text{and} \qquad \widehat{\mathcal{W}}_{\text{even}} = \widehat{W}^{J_{m-1}} \cdots \widehat{W}^{J_4}\widehat{W}^{J_2}\ .$$

Let $r' = \widehat{\mathcal{W}}_{\text{odd}}r$ and $r'' = \widehat{\mathcal{W}}_{\text{even}}r'$. Recall that for any vector s, $\widehat{W}^{J_k}s$ differs from s only in components indexed by J_k. Thus, by (B-1), (B-2), $\widehat{\mathcal{W}}_{\text{odd}}r$ can be studied by considering each term $\widehat{W}^{J_k}r$, k odd, separately. We have for k odd and $i \in J_k$

$$\begin{aligned} [\widehat{W}^{J_k}r]_i &= (W^{J_k}_{i,\partial_- J_k} + \epsilon)r_{\partial_- J_k} + (W^{J_k}_{i,\partial_+ J_k} + \epsilon)r_{\partial_+ J_k} + \epsilon\sum_{j\in J_k} r_j \\ &= (W^{J_k}_{i,\partial_- J_k} + W^{J_k}_{i,\partial_+ J_k} + 2\epsilon)a + \epsilon\sum_{j\in J_k} r_j \\ &\leq (W^{J_k}_{i,\partial_- J_k} + W^{J_k}_{i,\partial_+ J_k} + 2\epsilon)a + \epsilon|J_k|(a' \vee b)\ . \end{aligned}$$

[2] This is most easily seen by noting that the PAR sampler is equivalent to a two-block sampler, comprising the composite blocks: $J_{\text{odd}} := \bigcup_{k \text{ odd}} J_k$ and $J_{\text{even}} := \bigcup_{k \text{ even}} J_k$. The statement follows by applying Lemma 14 to these two blocks.

The second line follows by the assumption on r and the fact that boundaries of an odd numbered block lie in the adjacent even blocks. For $i \in J_{-k}^c$, $W_{i,\partial_- J_k}^{J_k} + W_{i,\partial_+ J_k}^{J_k} \leq \lambda$ and since $|J_k| \leq L$,

$$[\widehat{W}^{J_k} r]_i \leq \lambda a + \epsilon \left(2a + L(a' \vee b)\right) \ .$$

For $i \in J_k \cap J_{-k}$, $W_{i,\partial_- J_k}^{J_k} + W_{i,\partial_+ J_k}^{J_k} \leq \beta$ and thus

$$[\widehat{W}^{J_k} r]_i \leq \beta a + \epsilon \left(2a + L(a' \vee b)\right) \ .$$

It follows that, $r' = \widehat{\mathcal{W}}_{\text{odd}} r$,

$$[r']_i \leq \begin{cases} a, & i \in J_{-k}^c, \ k \text{ even}, \\ \lambda a + \epsilon \left(2a + L(a' \vee b)\right), & i \in J_{-k}^c, \ k \text{ odd}, \\ \beta a + \epsilon \left(2a + L(a' \vee b)\right), & i \in J_k \cap J_{-k}, \ k \text{ even}, \end{cases}$$

Similarly, $r'' = \widehat{\mathcal{W}}_{\text{even}} r'$ can be studied by considering each term $\widehat{W}^{J_k} r'$, k even, separately. Let $c = 2a + L(a' \vee b)$. For k even and $i \in J_k$,

$$\begin{aligned} [\widehat{W}^{J_k} r']_i &= (W_{i,\partial_- J_k}^{J_k} + \epsilon) r'_{\partial_- J_k} + (W_{i,\partial_+ J_k}^{J_k} + \epsilon) r'_{\partial_+ J_k} + \epsilon \sum_{j \in J_k} r'_j \\ &\leq (W_{i,\partial_- J_k}^{J_k} + W_{i,\partial_+ J_k}^{J_k} + 2\epsilon)(\lambda a + \epsilon c) + \epsilon \sum_{j \in J_k} r'_j \\ &\leq (W_{i,\partial_- J_k}^{J_k} + W_{i,\partial_+ J_k}^{J_k} + 2\epsilon)(\lambda a + \epsilon c) + \epsilon |J_k| \left(a(1 \vee \beta) + \epsilon c\right) \ . \end{aligned}$$

For $i \in J_{-k}^c$, $W_{i,\partial_- J_k}^{J_k} + W_{i,\partial_+ J_k}^{J_k} \leq \lambda$ and thus,

$$[\widehat{W}^{J_k} r']_i \leq \lambda^2 a + \epsilon \left(\lambda c + 2\lambda a + 2\epsilon c + La(1 \vee \beta) + \epsilon Lc\right) \ .$$

For $i \in J_k \cap J_{-k}$, $W_{i,\partial_- J_k}^{J_k} + W_{i,\partial_+ J_k}^{J_k} \leq \beta$ and thus,

$$[\widehat{W}^{J_k} r']_i \leq \lambda\beta a + \epsilon \left(\beta c + 2\lambda a + 2\epsilon c + La(1 \vee \beta) + \epsilon Lc\right) \ .$$

In summary, $r'' = \widehat{\mathcal{W}}_{\text{even}} r'$,

$$[r'']_i \leq \begin{cases} \lambda^2 a + \epsilon \left(\lambda c + 2\lambda a + 2\epsilon c + La(1 \vee \beta) + \epsilon Lc\right), & i \in J_{-k}^c, \ k \text{ even}, \\ \lambda a + \epsilon \left(2a + L(a' \vee b)\right), & i \in J_{-k}^c, \ k \text{ odd}, \\ \lambda\beta a + \epsilon \left(\beta c + 2\lambda a + 2\epsilon c + La(1 \vee \beta) + \epsilon Lc\right), & i \in J_k \cap J_{-k}, \ k \text{ even}. \end{cases}$$

□

A.4 Proof of Theorem 10

The proof is inductive. For any vector r, observe that $\widehat{W}^J r$ differs from r only in components indexed by J. In particular, for any $i \in J_1$, $[\widehat{W}^{J_1} \mathbf{1}]_i = W_{i,\partial_+ J_1}^{J_1} + (|J_1| + 1)\epsilon$. Thus,

$$[\widehat{W}^{J_1} \mathbf{1}]_i \leq \begin{cases} \lambda + c\epsilon, & \text{for } i \in J_{-1}^c, \\ \beta + c\epsilon, & \text{for } i \in J_1 \cap J_2. \end{cases}$$

We separately bound the terms for $i \in J_{-1}^c$ and $i \in J_1 \cap J_2$ since $W^{J_1}_{i,\partial_+ J_1}$ can approach 1 as i approaches the boundary $\partial_+ J_1$ (recall $\widehat{W}^{J_1}_{i,\partial_+ J_1} = W^{J_1}_{i,\partial_+ J_1} + \epsilon$ for $i \in J_1$).

Now assume

$$[\widehat{W}^{J_{k-1}} \cdots \widehat{W}^{J_1}\mathbf{1}]_i \le \begin{cases} \lambda + c\epsilon, & \text{for } i \in J^c_{-(k-1)} \text{ (middle of block)}, \\ \lambda' + 2c\epsilon, & \text{for } i \in J_{k-2} \cap J_{k-1} \text{ (left overlap)}, \\ \beta(\lambda \vee 1) + c\epsilon, & \text{for } i \in J_{k-1} \cap J_k \text{ (right overlap)}. \end{cases}$$

We have

$$\begin{aligned} [\widehat{W}^{J_k}\widehat{W}^{J_{k-1}} \cdots \widehat{W}^{J_1}\mathbf{1}]_i &= \sum_{j \in J_k^+} \widehat{W}^{J_k}_{i,j} [\widehat{W}^{J_{k-1}} \cdots \widehat{W}^{J_1}\mathbf{1}]_j \\ &= \sum_{j \in \partial J_k} (W^{J_k}_{i,j} + \epsilon)[\widehat{W}^{J_{k-1}} \cdots \widehat{W}^{J_1}\mathbf{1}]_j + \sum_{j \in J_k} \epsilon [\widehat{W}^{J_{k-1}} \cdots \widehat{W}^{J_1}\mathbf{1}]_j \ . \end{aligned} \tag{34}$$

The first term of the sum can be simplified to

$$\begin{aligned} &\sum_{j \in \partial J_k} (W^{J_k}_{i,j} + \epsilon)[\widehat{W}^{J_{k-1}} \cdots \widehat{W}^{J_1}\mathbf{1}]_j \\ &\qquad = (W^{J_k}_{i,\partial_- J_k} + \epsilon)[\widehat{W}^{J_{k-1}} \cdots \widehat{W}^{J_1}\mathbf{1}]_{\partial_- J_k} + (W^{J_k}_{i,\partial_+ J_k} + \epsilon)[\widehat{W}^{J_{k-1}} \cdots \widehat{W}^{J_1}\mathbf{1}]_{\partial_+ J_k} \\ &\qquad = (W^{J_k}_{i,\partial_- J_k} + \epsilon)(\lambda + c\epsilon) + (W^{J_k}_{i,\partial_+ J_k} + \epsilon) \\ &\qquad = W^{J_k}_{i,\partial_- J_k}\lambda + W^{J_k}_{i,\partial_+ J_k} + \epsilon(W^{J_k}_{i,\partial_- J_k} c + \lambda + c\epsilon + 1) \ . \end{aligned}$$

The second term of (34) is

$$\begin{aligned} \sum_{j \in J_k} \epsilon [\widehat{W}^{J_{k-1}} \cdots \widehat{W}^{J_1}\mathbf{1}]_j &= \sum_{j \in J_{k-1} \cap J_k} \epsilon [\widehat{W}^{J_{k-1}} \cdots \widehat{W}^{J_1}\mathbf{1}]_j + \sum_{j \in J^c_{k-1} \cap J_k} \epsilon \\ &\le |J_{k-1} \cap J_k| \epsilon(\beta(\lambda \vee 1) + c\epsilon) + \epsilon |J^c_{k-1} \cap J_k| \\ &\le \epsilon(\beta(\lambda \vee 1) \vee 1)|J_k| + c\epsilon^2 |J_{k-1} \cap J_k| \ . \end{aligned}$$

Thus

$$\begin{aligned} [\widehat{W}^{J_k}\widehat{W}^{J_{k-1}} \cdots \widehat{W}^{J_1}\mathbf{1}]_i \le W^{J_k}_{i,\partial_- J_k}\lambda + W^{J_k}_{i,\partial_+ J_k} \\ + \epsilon(W^{J_k}_{i,\partial_- J_k} c + \lambda + c\epsilon + 1 + (\beta(\lambda \vee 1) \vee 1)|J_k| + c\epsilon |J_{k-1} \cap J_k|) \ . \end{aligned}$$

For $i \in J_k \cap J_{k+1}$, $W^{J_k}_{i,\partial_- J_k}\lambda + W^{J_k}_{i,\partial_+ J_k} \le \beta(\lambda \vee 1)$. For $i \in J^c_{-k}$, $W^{J_k}_{i,\partial_- J_k}\lambda + W^{J_k}_{i,\partial_+ J_k} \le \lambda$ by the definition of λ. For $i \in J_{k-1} \cap J_k$, $W^{J_k}_{i,\partial_- J_k}\lambda + W^{J_k}_{i,\partial_+ J_k} \le \lambda'$ by the definition of λ'.

For $i \in J_k \cap J^c_{k-1}$, the coefficient of ϵ is by itself bounded by c since

$$\frac{\lambda + 1 + (\beta(\lambda \vee 1) \vee 1)|J_k|}{1 - W^{J_k}_{i,\partial_- J_k} - \epsilon - \epsilon |J_{k-1} \cap J_k|} \le c \ .$$

For $i \in J_{k-1} \cap J_k$, the coefficient of ϵ is

$$\begin{aligned} W^{J_k}_{i,\partial_- J_k} c + \lambda + c\epsilon + 1 + (\beta(\lambda \vee 1) \vee 1)|J_k| + c\epsilon |J_{k-1} \cap J_k| \\ \le c + \lambda + c\epsilon + 1 + (\beta(\lambda \vee 1) \vee 1)|J_k| + c\epsilon |J_{k-1} \cap J_k|) \le 2c \ . \end{aligned}$$

□

B Additional Proofs

B.1 Proof of Lemma 3

Lemma 16. *For each $i \in 1, \ldots, t$, let $P_i(x_{i-1}, \mathrm{d}x_i)$ be a Markov transition kernel on X. For some integer $h > 0$, assume that the composite transition kernels $Q_i(x_{(i-1)h}, \mathrm{d}x_{ih}) = (P_{(i-1)h+1} \ldots P_{ih})(x_{(i-1)h}, \mathrm{d}x_{ih})$ satisfy the following minorisation condition: there exists probability measures $\nu_i(\mathrm{d}x)$ on X and a common constant $\alpha \in [0,1]$ such that $Q_i \geq \alpha\nu_i$. Then for the probability measures (on the product space X^t) $\prod_{i=1}^t P(x_{i-1}, \mathrm{d}x_i)$ and $\prod_{i=1}^t P(y_{i-1}, \mathrm{d}y_i)$, and for any x_0 and y_0, there exists a coupling Ψ such that if $(X_{1:t}, Y_{1:t}) \sim \Psi$ then $\Pr(X_i \neq Y_i) \leq (1-\alpha)^{\lfloor i/h \rfloor}$.*

Proof. When $h = 1$ the results follows standard arguments (Lindvall, 2002) and is repeated here for the sake of completeness. The coupling Ψ attempts to couple (X_1, Y_1), followed by (X_2, Y_2) etc. Specifically, if $X_{i-1} = Y_{i-1}$, then draw X_i from $P_i(X_{i-1}, \mathrm{d}x_i)$ and set $Y_i = X_i$. If $X_{i-1} \neq Y_{i-1}$, draw (X_i, Y_i) from the measure

$$\alpha\nu_i(\mathrm{d}x_i)\delta_{x_i}(\mathrm{d}y_i) + (1-\alpha)^{-1}\left(P_i(X_{i-1}, \mathrm{d}x_i) - \alpha\nu(\mathrm{d}x_i)\right)\left(P_i(Y_{i-1}, \mathrm{d}y_i) - \alpha\nu(\mathrm{d}y_i)\right)$$

where δ_x is the atom measure. It now follows that $\Pr(X_i \neq Y_i) \leq (1-\alpha)\Pr(X_{i-1} \neq Y_{i-1}) \leq (1-\alpha)^i$.

For $h > 1$ we couple the skeleton process $(X_h, Y_h), (X_{2h}, Y_{2h}), \ldots$ using a similar scheme as for $h = 1$. (Note that the h-skeleton X-system has transition kernels Q_i that satisfy a minorisation condition which is to be used in the same manner as in the proof for $h = 1$.) Let i be the first instance that $X_{ih} = Y_{ih}$, i.e. $X_{jh} \neq Y_{jh}$ for $j < i$. Then, simulate the future X-process, X_k for $k = ih+1, ih+2, \ldots$ from P_k and set $Y_k = X_k$. The non-skeleton terms X_k and Y_k for $k < ih$ are simulated independently from their respective conditional laws. Thus for $i \geq h$, $\Pr(X_i \neq Y_i) \leq \Pr(X_{h\lfloor i/h \rfloor} \neq Y_{h\lfloor i/h \rfloor})$ which is bounded above by $(1-\alpha)^{\lfloor i/h \rfloor}$. □

Proof. (Lemma 3) For notational brevity we write $\mathrm{d}x$ instead of $\nu(\mathrm{d}x)$, where ν is the dominating probability measure defined in (S-1).

Recall that under (B-1), J is an interval and we can write $J = \{s, \ldots, u\}$. To compute $W_{i,j}^J$ for $i \in J$ we use a coupling as in (18):

$$W_{i,j}^J := \sup_{\substack{x,\, z \,\in\, \mathsf{X}^n \\ x_{-j} \,=\, z_{-j}}} \Psi_{j,x,z}^J(X_i' \neq Z_i'), \tag{35}$$

where $\Psi_{j,x,z}^J$ is a coupling of ϕ_x^J and ϕ_z^J for $x_{-j} = z_{-j}$. We know from Lemma 2 that we only need to consider the cases $j = s-1$ and $j = u+1$. Consider first $j = s-1$ (assuming $s > 1$).

Write the density of ϕ_x^J as $p(x_s, \ldots, x_u | x_{s-1}, x_{u+1}, y_{s:u}) = \prod_{i=s}^u p(x_i | x_{i-1}, x_{u+1}, y_{i:u})$, which is a product of inhomogeneous Markov transition kernels. (A similar expression follows if written backwards, i.e. $\prod_{i=s}^u p(x_i | x_{i+1}, x_{s-1}, y_{s:i})$.)

We show that the composite kernel formed by any h (for h defined in (S-1) and (S-2)) successive kernels of the given inhomogeneous product, i.e. $\int \prod_{i=t+1}^{t+h} p(x_i | x_{i-1}, x_{u+1}, y_{i:u}) \mathrm{d}x_{t+1:t+h-1}$, satisfies a minorisation condition with respect to some probability measure and the time-uniform constant $\delta^{\frac{1-h}{h}} \frac{\sigma_-}{\sigma_+}$. Thus the coupling $\Psi_{j,x,z}^J$ can be defined as in Lemma 16 to complete the proof.

Let $t \leq u - h$ be some time index and consider

$$p(x_{t+h} \,|\, x_t, x_{u+1}, y_{t+1:u}) = \frac{p(x_{u+1}, y_{t+h:u} \,|\, x_{t+h}) p(x_{t+h} \,|\, x_t, y_{t+1:t+h-1})}{\int p(x_{u+1}, y_{t+h:u} \,|\, x_{t+h}) p(x_{t+h} \,|\, x_t, y_{t+1:t+h-1}) \mathrm{d}x_{t+h}}. \tag{36}$$

We have

$$p(x_{t+h} \,|\, x_t, y_{t+1:t+h-1}) = \frac{\int \prod_{j=t+1}^{t+h-1} \{g(x_j, y_j) m(x_{j-1}, x_j)\}\, m(x_{t+h-1}, x_{t+h}) \mathrm{d}x_{t+1:t+h-1}}{\int \prod_{j=t+1}^{t+h-1} \{g(x_j, y_j) m(x_{j-1}, x_j)\}\, m(x_{t+h-1}, x_{t+h}) \mathrm{d}x_{t+1:t+h}},$$

which implies, using (S-1) and (S-2),

$$\delta^{\frac{1-h}{h}} \sigma_- \leq p(x_{t+h} \,|\, x_t, y_{t+1:t+h-1}) \leq \sigma_+.$$

Plugging these two bounds into the numerator and denominator, respectively, of (36) we get

$$p(x_{t+h} \,|\, x_t, x_{u+1}, y_{t+1:u}) \geq \delta^{\frac{1-h}{h}} \frac{\sigma_-}{\sigma_+} \frac{p(x_{u+1}, y_{t+h:u} \,|\, x_{t+h})}{\int p(x_{u+1}, y_{t+h:u} \,|\, x_{t+h}) \mathrm{d}x_{t+h}}. \tag{37}$$

The proof for $j = u + 1$ follows analogously by considering a backward decomposition of $p(x_i \,|\, x_{s-1}, x_{u+1}, y_{s:u})$ and minorising $p(x_{t-h} \,|\, x_t, x_{s-1}, y_{s:t-1})$. □

B.2 Proof of Proposition 12

B.2.1 Part 1 – Uniform minorisation of the blocked PG kernel

We consider the blocked PG kernel for block $J = \{s, \ldots, u\}$, Q_N^J. The following proof is a slight modification of the proof of Proposition 5 by Lindsten et al. (2015), which amount to taking the conditioning on the end-point x_{u+1} into account under the h-step forgetting condition (S-1)b. It is known from Lindsten et al. (2015, Theorem 1) that, for any $x_{J^+} \in \mathsf{X}^{|J|+2}$ and measurable $A \subset \mathsf{X}^{|J|}$, $Q_N^J(x_{J^+}, A) \geq (1 - \epsilon(N, |J|))\phi_x^J(A)$, where

$$\epsilon(N, |J|) = 1 - \prod_{t=s}^{u} \frac{N-1}{2B_t^J + N - 2},$$

$$B_t^J = \left\{ \sup_{x_t} \frac{p(y_{t:u}, x_{u+1} \,|\, x_t)}{p(y_{t:u}, x_{u+1} \,|\, x_{s-1}, y_{s:t-1})} \right\} \vee \max_{0 \leq \ell < u-t} \left\{ \sup_{x_t} \frac{p(y_{t:t+\ell} \,|\, x_t)}{p(y_{t:t+\ell} \,|\, x_{s-1}, y_{s:t-1})} \right\}.$$

Assume (S-1),(S-1). From Lindsten et al. (2015, Proposition 5) if follows that the second term in the definition of B_t^J is bounded by $\delta \frac{\sigma_+}{\sigma_-}$. It remains to bound the first term, incorporating the dependence on the boundary point x_{u+1}.

Consider first $t \leq u - h + 1$. We have for the numerator

$$\begin{aligned}
p(y_{t:u}, x_{u+1} \,|\, x_t) &= \int \prod_{j=t}^{t+h-1} (g(x_j, y_j) m(x_j, x_{j+1})) \prod_{j=t+h}^{u} (g(x_j, y_j) m(x_j, x_{j+1}))\, \mathrm{d}x_{t+1:u} \\
&\leq \prod_{j=t}^{t+h-1} \left\{ \sup_x g(x, y_j) \right\} \int \prod_{j=t}^{t+h-1} m(x_j, x_{j+1}) \prod_{j=t+h}^{u} (g(x_j, y_j) m(x_j, x_{j+1}))\, \mathrm{d}x_{t+1:u} \\
&\leq \sigma_+ \prod_{j=t}^{t+h-1} \left\{ \sup_x g(x, y_j) \right\} \int \prod_{j=t+h}^{u} (g(x_j, y_j) m(x_j, x_{j+1}))\, \mathrm{d}x_{t+h:u}\,.
\end{aligned}$$

Analogously we get for the denominator

$$\begin{aligned}
p(y_{t:u}, x_{u+1} \,|\, x_{s-1}, y_{s:t-1}) &= \int p(y_{t:u}, x_{u+1} \,|\, x_t) p(x_t \,|\, x_{s-1}, y_{s:t-1}) \mathrm{d}x_t \\
&\geq \sigma_- \prod_{j=t}^{t+h-1} \left\{ \inf_x g(x, y_j) \right\} \int \prod_{j=t+h}^{u} (g(x_j, y_j) m(x_j, x_{j+1}))\, \mathrm{d}x_{t+h:u}\,.
\end{aligned}$$

It follows that for $t \leq u - h + 1$,

$$\sup_{x_t} \frac{p(y_{t:u}, x_{u+1} \,|\, x_t)}{p(y_{t:u}, x_{u+1} \,|\, x_{s-1}, y_{s:t-1})} \leq \delta \frac{\sigma_+}{\sigma_-}$$

Next, consider the case $t > u - h + 1$. We have for the numerator

$$\begin{aligned} p(y_{t:u}, x_{u+1} \,|\, x_t) &= \int \prod_{j=t}^{u} \left(g(x_j, y_j) m(x_j, x_{j+1})\right) \mathrm{d}x_{t+1:u} \\ &\leq \prod_{j=t}^{u} \left\{ \sup_x g(x, y_j) \right\} \int \prod_{j=t}^{u} m(x_j, x_{j+1}) \mathrm{d}x_{t+1:u} \leq \sigma_+ \prod_{j=t}^{u} \left\{ \sup_x g(x, y_j) \right\} . \end{aligned}$$

For the denominator we write

$$p(y_{t:u}, x_{u+1} \,|\, x_{s-1}, y_{s:t-1}) = \frac{p(y_{u-h+1:u}, x_{u+1} \,|\, x_{s-1}, y_{s:u-h})}{p(y_{u-h+1:t-1} \,|\, x_{s-1}, y_{s:u-h})} .$$

We have

$$\begin{aligned} &p(y_{u-h+1:u}, x_{u+1} \,|\, x_{s-1}, y_{s:u-h}) \\ &\qquad = \int \prod_{j=u-h+1}^{u} \left(g(x_j, y_j) m(x_j, x_{j+1})\right) p(x_{u-h+1} \,|\, x_{s-1}, y_{s:u-h}) \mathrm{d}x_{u-h+1:u} \\ &\quad \geq \prod_{j=u-h+1}^{u} \left\{ \inf_x g(x, y_j) \right\} \int \prod_{j=u-h+1}^{u} m(x_j, x_{j+1}) p(x_{u-h+1} \,|\, x_{s-1}, y_{s:u-h}) \mathrm{d}x_{u-h+1:u} \\ &\qquad \geq \sigma_- \prod_{j=u-h+1}^{u} \left\{ \inf_x g(x, y_j) \right\} , \end{aligned}$$

and

$$\begin{aligned} &p(y_{u-h+1:t-1} \,|\, x_{s-1}, y_{s:u-h}) \\ &\qquad = \int \prod_{j=u-h+1}^{t-1} \left(g(x_j, y_j) m(x_j, x_{j+1})\right) p(x_{u-h+1} \,|\, x_{s-1}, y_{s:u-h}) \mathrm{d}x_{u-h+1:t} \\ &\qquad \leq \sigma_+ \prod_{j=u-h+1}^{t-1} \left\{ \sup_x g(x, y_j) \right\} . \end{aligned}$$

In summary we get

$$\sup_{x_t} \frac{p(y_{t:u}, x_{u+1} \,|\, x_t)}{p(y_{t:u}, x_{u+1} \,|\, x_{s-1}, y_{s:t-1})} \leq \delta \frac{\sigma_+}{\sigma_-}$$

also for $t > u - h + 1$, and thus $B_t^J \leq \delta \frac{\sigma_+}{\sigma_-}$ for all $t \in J$. The result follows. □

B.2.2 Part 2 – Wasserstein estimate for the blocked PG kernel

The statement is a corollary to the following lemma.

Lemma 17. *Let P and Q be two Markov kernels on X^n and assume that there exists a constant $\epsilon \in [0,1]$ such that $Q(x, \mathrm{d}y) \geq (1-\epsilon)P(x, \mathrm{d}y)$ for all $x \in \mathsf{X}^n$. For $x, y \in \mathsf{X}^n$ such that $x_{-j} = y_{-j}$, let $\Psi_{j,x,y}$ be a coupling of $P(x, \cdot)$ and $P(y, \cdot)$ and let W be the matrix (a Wasserstein matrix for P)*

$$W_{i,j} = \sup_{\substack{x,\, y \in \mathsf{X}^n \\ x_{-j} = y_{-j}}} \Psi_{j,x,y}(X_i' \neq Y_i') = \sup_{\substack{x,\, y \in \mathsf{X}^n \\ x_{-j} = y_{-j}}} \int \Psi_{j,x,y}(\mathrm{d}x', \mathrm{d}y') \mathbb{I}_{[x_i' \neq y_i']}.$$

Then, the matrix $\widehat{W}$ with $\widehat{W}_{i,j} = W_{i,j} + \epsilon$ is a Wasserstein matrix for Q.

Proof. For $\epsilon = 0$ the result is immediate. Hence, consider $\epsilon > 0$. For any $x, y \in \mathsf{X}^n$, define $r_x = \frac{1}{\epsilon}(Q(x, \cdot) - (1-\epsilon)P(x, \cdot))$ and let $R_{x,y}$ be a coupling of r_x and r_y. It follows that, for any $x, y \in \mathsf{X}^n$ with $x_{-j} = y_{-j}$, $(1-\epsilon)\Psi_{j,x,y} + \epsilon R_{x,y}$ is a coupling of $Q(x, \cdot)$ and $Q(y, \cdot)$. We have,

$$(1-\epsilon)\Psi_{j,x,y}(X_i' \neq Y_i') + \epsilon R_{x,y}(X_i' \neq Y_i') \leq (1-\epsilon)W_{i,j} + \epsilon \leq \widehat{W}_{i,j} \ .$$

□

References

Andrieu, C., A. Doucet, and R. Holenstein (2010). Particle Markov chain Monte Carlo methods. *Journal of the Royal Statistical Society: Series B 72*(3), 269–342.

Andrieu, C., A. Lee, and M. Vihola (2015, April). Uniform ergodicity of the iterated conditional SMC and geometric ergodicity of particle Gibbs samplers. arXiv.org, arXiv:1312.6432v2.

Cappé, O., E. Moulines, and T. Rydén (2005). *Inference in Hidden Markov Models*. Springer.

Carter, C. K. and R. Kohn (1994). On Gibbs sampling for state space models. *Biometrika 81*(3), 541–553.

Chopin, N. and S. S. Singh (2015). On particle Gibbs sampling. *Bernoulli 21*(3), 1855–1883.

Del Moral, P. (2004). *Feynman-Kac Formulae - Genealogical and Interacting Particle Systems with Applications*. Probability and its Applications. Springer.

Doucet, A., S. J. Godsill, and C. Andrieu (2000). On sequential Monte Carlo sampling methods for Bayesian filtering. *Statistics and Computing 10*(3), 197–208.

Doucet, A. and A. Johansen (2011). A tutorial on particle filtering and smoothing: Fifteen years later. In D. Crisan and B. Rozovskii (Eds.), *The Oxford Handbook of Nonlinear Filtering*. Oxford University Press.

Follmer, H. (1982). A covariance estimate for Gibbs measures. *Journal of Functional Analysis 46*, 387–395.

Frühwirth-Schnatter, S. (1994). Data augmentation and dynamic linear models. *Journal of Time Series Analysis 15*(2), 183–202.

Kuhlenschmidt, B. (2014). *On the stability of Sequential Monte Carlo methods for parameter estimation*. Ph. D. thesis, Cambridge Centre for Analysis, DPMMS, Cambridge University.

Lindsten, F., R. Douc, and E. Moulines (2015). Uniform ergodicity of the particle Gibbs sampler. *Scandinavian Journal of Statistics 42*(3), 775–797.

Lindsten, F., M. I. Jordan, and T. B. Schön (2014). Particle Gibbs with ancestor sampling. *Journal of Machine Learning Research 15*, 2145–2184.

Lindsten, F. and T. B. Schön (2013). Backward simulation methods for Monte Carlo statistical inference. *Foundations and Trends in Machine Learning 6*(1), 1–143.

Lindvall, T. (2002). *Lectures on the Coupling Method*. Dover.

Rebeschini, P. and R. van Handel (2014). Comparison theorems for Gibbs measures. *Journal of Statistical Physics 157*(2), 234–281.

Wang, N.-Y. and L. Wu (2014). Convergence rate and concentration inequalities for Gibbs sampling in high dimensions. *Bernoulli 20*(4), 1698–1716.

Whiteley, N. (2010). Discussion on Particle Markov chain Monte Carlo methods. *Journal of the Royal Statistical Society: Series B 72*(3), 306–307.

Whiteley, N., C. Andrieu, and A. Doucet (2010). Efficient Bayesian inference for switching state-space models using discrete particle Markov chain Monte Carlo methods. Technical report, Bristol Statistics Research Report 10:04.